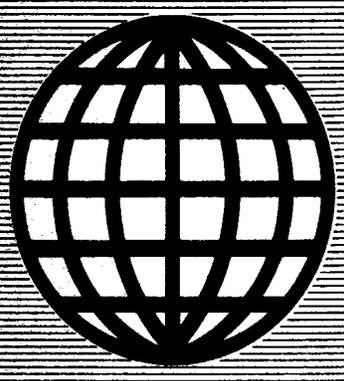

The COMPUTER JOURNAL

Programming - User Support

Applications

Issue Number 58 November/December 1992 US$3.95

Z-System Corner

Dr. S-100

Development Tools

Real Computing

Kaypro Reset Move

Computing Timer Values

Multitasking in Forth

The Computer Corner

ISSN # 0748-9331

EPROM PROGRAMMERS

Stand-Alone Gang Programmer $750.00

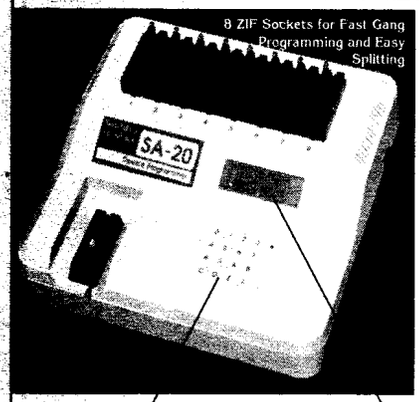

- Completely stand-alone or PC driven
- Programs E(E)PROMs
- ***1 Megabit of DRAM***
- ***User upgradable to 32 Megabit***
- ***.3/.6" ZIF socket, RS-232, Parallel In and Out***
- 32K internal Flash EEPROM for easy firmware upgrades
- ***Quick Pulse Algorithm (27256 in 5 sec, 1 Megabit in 17 sec.)***
- 2 year warranty
- Made in U.S.A.
- Technical support by phone
- Complete manual and schematic
- ***Single Socket Programmer also available. $550.00***
- Split and Shuffle 16 & 32 bit
- 100 User Definable Macros, 10 User Definable Configurations
- Intelligent Identifier
- Binary, Intel Hex, and Motorola S

20 Key Tactile Keypad (not membrane) **20 x 4 Line LCD Display**

Internal Programmer for PC $139.95

New Intelligent Averaging Algorithm. Programs 64A in 10 sec., 256 in 1 min., 1 Meg (27010, 011) in 2 min. 45 sec., 2 Meg (27C2001) in 5 min. Internal card with external 40 pin ZIF.

2 ft. Cable **40 pin ZIF**

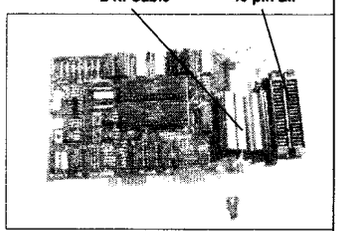

- Reads, verifies, and programs 2716, 32, 32A, 64, 64A, 128, 128A, 256, 512, 513, 010, 011, 301, 27C2001, MCM 68764, 2532
- ***Automatically sets programming voltage***
- Load and save buffer to disk
- Binary, Intel Hex, and Motorola S formats
- ***Upgradable to 32 Meg EPROMs***
- ***No personality modules required***
- 1 year warranty • 10 day money back guarantee
- Adapters available for 8748, 49, 51, 751, 52, 55, TMS 7742, 27210, 57C1024, and memory cards
- Made in U.S.A.

NEEDHAM'S ELECTRONICS
4539 Orange Grove Ave. • Sacramento, CA 95841
Mon. - Fri. 8am - 5pm PST

C.O.D.

Call for more information
(916) 924-8037
FAX (916) 972-9960

Cross-Assemblers as low as $50.00
Simulators as low as $100.00
Cross-Disassemblers as low as $100.00
Developer Packages
as low as $200.00(a $50.00 Savings)

A New Project
Our line of macro Cross-assemblers are easy to use and full featured, including conditional assembly and unlimited include files.

Get It To Market--FAST
Don't wait until the hardware is finished to debug your software. Our Simulators can test your program logic before the hardware is built.

No Source!
A minor glitch has shown up in the firmware, and you can't find the original source program. Our line of disassemblers can help you re-create the original assembly language source.

Set To Go
Buy our developer package and the next time your boss says "Get to work.", you'll be ready for anything.

Quality Solutions
PseudoCorp has been providing quality solutions for microprocessor problems since 1985.

BROAD RANGE OF SUPPORT
- Currently we support the following microprocessor families (with more in development):

Intel 8048	RCA 1802,05	Intel 8051	Intel 8096
Motorola 6800	Motorola 6801	Motorola 68HC11	Motorola 6805
Hitachi 6301	Motorola 6809	MOS Tech 6502	WDC 65C02
Rockwell 65C02	Intel 8080,85	Zilog Z80	NSC 800
Hitachi HD64180	Motorola 68000,8	Motorola 68010	Intel 80C196

- All products require an IBM PC or compatible.

So What Are You Waiting For? Call us:

PseudoCorp
Professional Development Products Group
716 Thimble Shoals Blvd, Suite E
Newport News, VA 23606
(804) 873-1947 **FAX: (804)873-2154**

FORTH

Journey with us to discover the shortest path between programming problems and efficient solutions.

The Forth programming language is a model of simplicity: In about 16K, it can offer a complete development system in terms of compiler, editor, and assembler, as well as an interpretive mode to enhance debugging, profiling, and tracing.

As an "open" language, Forth lets you build new control-flow structures, and other compiler-oriented extensions that closed languages do not.

Forth Dimensions is the magazine to help you along this journey. It is one of the benefits you receive as a member of the non-profit Forth Interest Group (FIG). Local chapters, the GEnie™ Forth RoundTable, and annual FORML conferences are also supported by FIG. To receive a mail-order catalog of Forth literature and disks, call 510-89-FORTH or write to:
Forth Interest Group, P.O. Box 2154, Oakland, CA 94621.
Membership dues begin at $40 for the U.S.A. and Canada.
Student rates begin at $18 (with valid student I.D.).

GEnie is a trademark of General Electric.

SAGE MICROSYSTEMS EAST

Selling and Supporting the Best in 8-Bit Software

Z3PLUS or NZCOM (now only $49 each)
XBIOS for SB180 ($50)
ZMATE text editor ($50)
BDS C for Z-system (only $60)
DSD: Dynamic Screen Debugger ($50)
PCED: ARUNZ and LSH for MSDOS ($50)
ZMAC macro-assembler ($50, $70 with printed manual)
Order by phone, mail, or modem and use
Check, VISA, or MasterCard.

Z-System public domain software by mail.

Regular Subscription Service
Z3COM Package of over 1.5 MB of COM files
Z3HELP Package with over 1.3 MB of online documentation
Z-SUS Programmers Pack, 8 disks full
Z-SUS Word Processing Toolkit
And More!

For catalog on disk, send $2.00 ($4.00 outside North America) and your computer format to:

Sage Microsystems East
1435 Centre Street
Newton Centre MA 02159-2469
(617) 965-3552 (voice 9 to 11AM)
(617) 965-7259 (pw=DDT)
(MABOS on PC-Pursuit)

The Computer Journal

Founder
Art Carlson

Editor/Publisher
Bill D. Kibler

Technical Consultant
Chris McEwen

Contributing Editors
Brad Rodriquez
Matt Mercaldo
Tim McDonough
Frank Sergeant
Clem Pepper
Richard Rodman
Jay Sage

The Computer Journal is published six times a year and mailed from *The Computer Journal*, P. O. Box 535, Lincoln, CA 95648, (916) 645-1670.

Opinions expressed in *The Computer Journal* are those of the respective authors and do not necessarily reflect those of the editorial staff or publisher.

Entire contents copyright © 1992 by *The Computer Journal* and respective authors. All rights reserved. Reproduction in any form prohibited without express written permission of the publisher.

Subscription rates within the US: $18 one year (6 issues), $32 two years (12 issues). Foreign (surface rate): $24 one year, $44 two years. Foreign (airmail): $38 one year, $72 two years. All funds must be in U.S. dollars drawn on a U.S. bank.

Send subscription, renewals, address changes, or advertising inquires to: *The Computer Journal*, P.O. Box 535, Lincoln,CA 95648.

Registered Trademarks

It is easy to get in the habit of using company trademarks as generic terms, but these trademarks are the property of the respective companies. It is important to acknowledge these trademarks as their property to avoid their losing the rights and the term becoming public property. The following frequently used trademarks are acknowledged, and we apologize for any we have overlooked.

Apple II, II+, IIc, IIe, Lisa, Macintosh, ProDos; Apple Computer Company. CP/M, DDT, ASM, STAT, PIP; Digital Research. DateStamper, BackGrounder ii, Dos Disk; Plu*Perfect Systems. Clipper, Nantucket; Nantucket, Inc. dBase, dBASE II, dBASE III, dBASE III Plus, dBASE IV; Ashton-Tate, Inc. MBASIC, MS-DOS, Windows, Word; MicroSoft. WordStar; MicroPro International. IBM-PC, XT, and AT, PC-DOS; IBM Corporation. Z80, Z280; Zilog Corporation. Turbo Pascal, Turbo C, Paradox; Borland International. HD64180; Hitachi America, Ltd. SB180; Micromint, Inc.

Where these and other terms are used in The Computer Journal, they are acknowledged to be the property of the respective companies even if not specifically acknowledged in each occurrence.

TCJ The Computer Journal
Issue Number 58 November / December 1992

Editor's Comments 2

Reader to Reader 3

Z-Systems Corner 7
Language Independence, part two.
By Jay Sage.

Real Computing 11
Minix, UZI, and GNU.
By Rick Rodman.

Affordable Development Tools 15
Finding affordable microcomputer development tools.
By Tim McDonough.

DR. S-100 17
Resurrecting your S-100 system.
By Herb R. Johnson.

Mr. Kaypro 20
Moving the reset botton.
By Charles B. Stafford.

Computing Timer Values. 21
Finding monostable values using "C".
By Clem Pepper.

Multitasking Forth 28
Forth Multitasking in a nutshell.
By Brad Rodriguez.

The Computer Corner 40
By Bill Kibler.

EDITOR'S COMMENTS

Welcome to the holidays. We don't have any special seasonal articles for you in this issue. However I did get a welcome gift as I was finishing up the magazine. Dave Thompson the owner of MicroCornucopia sent me all his Kaypro disk. Now TCJ is ready to sell you those Kaypro disk that you have been waiting to get. I am printing half the list in this issue, with the second part next time. That list is on the last few pages of the magazine.

Our writers have been busy as usual producing some great material for you. We have Jay Sage continuing his discussion on Language Independence. Tim McDonough is back with a report on microprocessor development tools. Rick Rodman comments on Minix and UZI (a Z80 UNIX?). We start out Kaypro support with Mr. Kaypro by Charles Stafford. Herbert Johnson explains how to get started bringing up an S-100 system.

We have two special articles for you, one to bring back those fundamentals of electronics, and one on multitasking in Forth. Clem Pepper show you a great use for "C", by explaining and giving you the source code for a simple program to calculate the resistors and capacitors needed for timers and monostable oscillators.

Brad Rodriguez tells us all and more about the Forth multitasker. This in depth study should answer questions about Forth's multitasker, as well as giving examples and alternatives to the way F83 works now.

Doing Business

I have been working very hard to get TCJ's production problems under control. The subscription problems however continue. It is becoming apparent that a new mailing label program is needed. I have been using the old program, but far too many items are not handled correctly. For the reader this means more time until mailing related problems go away. Please bear with me as I try and correct this problem.

When phoning TCJ please remember that my office hours are 9 to 11 PM week nights in California. My employer has dropped the alternate Fridays off, so I no longer can make daytime phone calls except on weekends. This makes doing business very difficult, but there are no alternatives at the present.

I will not be teaching next year and so two more nights a week will be available to catch up on mail. Currently I only answer mail and phone messages by letter between issue preparations. It goes something like this: edit stories non-stop; page set the stories in a rush, do a panic collate the material into a magazine (ignoring errors in hopes of being only a week or two late and not two months behind); take to printer quickly; update last months worth of phone calls and subscription updates; print labels before printer has issues ready (takes 3 hours to print if no errors happen); get issues back and put labels on over two very long nights; take time off from work to mail out the issues; relax for one day; start contacting writers; do mail; start again.

To add to all this frantic chaos, I must dig into my own pocket each issue to cover expenses beyond subscription returns. I have calculated cost at $2.75 each to mail in the US. So far renewals and new subscriptions have not been keeping pace with expenses (about 30% short). The results of this will be subscription increases next year, about a dollar per issue for now. This is far short of what some have suggested I charge ($50 a year), but well within what others are charging for like publications. As always I am open to suggestions and alternative proposals.

Till next issue, read, enjoy. and be merry, Bill Kibler.

READER to READER

Letters to the Editor

All Readers

MINI Articles

8/11/92

Dear Mr. Kibler

I just received the extra "free" copy of TCJ, and was fairly pleased at the changes discussed.

My perspective on things:

I am one of those that has been looking for a good alternative to the Byte/Creative Computing of yesteryear. Today's Byte is not much more than PC Magazine with a few features on bigger personal computer systems.

Dr. Dobbs Journal often has interesting material, but has headed very much into software engineering, which may be useful, but isn't much fun.

I saw one issue of MicroCornucopia, which seemed to be just about perfect in having coverage of various levels of "hobby computing." Even the articles that I didn't have any thoughts on implementing were highly interesting. They even had some people who could actually write, and an editor with a good sense of humor. I understand that that issue may have been the last one, alas!

Then I heard about TCJ, which was a little more hardware oriented than I am, and was foremost a CP/M and Z-System journal. But I heard that they were heading towards more support of Forth, and tried to do various sorts of coverage. Sounded pretty neat.

After a year, it seemed to be a little too oriented towards CP/M stuff, in which I have relatively little interest. Issue #56 seems to be better oriented, from my perspective.

I'll just quickly describe my computer system:
- Atari Mega STe, 4MB RAM
- SCSI Hard drive (85MB, mostly full!)
- Using GNU things like GCC (GNU C), and **perl** (Practical Extraction and Reporting Language OR Pathologically Eclectic Rubbish Lister ...Probably worthy of an article in TCJ.)
- Multitasking system called MiNT (MiNT is Not TOS)

People that want to do UNIX-like stuff on the ST should check out MiNT (I can send out copies - it uses the GNU Copyleft license, and is free!). It tends to crash quite a bit on me, which is partly a function of having a lot of system utilities that "break OS rules" and a function of working on fairly complex code that crashes far too much...It would work better with an MMU, I'm sure.

It's probably a more viable system than MINIX, particularly since it actually has some Atari support, and extensive developer support on Usenet.

- The LATEX typesetting system. Its fundamental purpose is to do "nasty math," like:

(editor comment:...It's does such a nice job of nasty display that I can't even try matching them...take my word it looked great in his letter...)

It produces VERY nice output, and is extremely portable.

- Many implementations of Forth for the Atari ST

I know of only two version that I don't have, one being the Forth-83 for the 68000, and the other being some f-83 commercial variant called MutliForth. I probably ought to do an article for TCJ on the various implementations that I do have. (Mitch Bradley's Forthmacs, The Australian FORST, H&D Forth, 4xForth) each certainly has advantages and disadvantages. 3 of these are commercial, so that people might be quite pleased if they got reviewed. (There would be some kind words for each one of the bunch.)

My current "big project" is some numerical analysis work, solving some fairly large scale optimization problems. More specifically, I am nearing completion of a Master's thesis on the topic: "Using Interior Point Methods to Solve the Multicommodity Network Flow Problem." When I'm not writing letters, I'm debugging code for some combination of linear programs and network flows.

One thing that is notable about all this stuff is that unlike the 8 bit users, I don't have any important applications that are NOT memory intensive. There isn't even one program (of importance) that takes less than 64K of RAM. The various compliers vary fairly wildly in memory consumption - GCC would not work when I had only 2MB, and I think it would be happier if I had more memory. Perl dynamically generates hash tables to build "association lists," and may take up a whole lot of RAM, or (relatively!) little if there's only a little data.

Something that I would like to see/get is some sort of small single-board computer that could act as an experimental controller or perhaps as a home-control coprocessor within my computer system. I've got this HUGE case in which my hard drive is mounted. There's power to spare, cables, and quite a lot of space. Every time I look around, I see another

different single board computer that wants to be used.

At this time, money is certainly a major restriction for me. The question is: Should I find some 8031/51 Kit, find some way of getting some old Z80 design to run, or maybe pick up one of those New Micros "Single Chip Computers?" There's no old code for me to support - I can go for something new without breaking anything. What would be nicest would be to have something running Forth, with an extra serial line to tether to either the VT100 or the ST. I'd like to see an article discussing some of the pros and cons of different small controllers.

I hope that this information is useful, and that it can help (to some extent) to guide where things go at TCJ. I'd certainly be willing and able to write an article or two, within the areas that I know. I've mentioned a couple of my ideas - if you have any thoughts, I wouldn't mind hearing them!

Yours truly, Christopher Browne, University of Ottawa

Well Mr. Browne:

I was unable to print your letter last issue, but your letter only seems to get better with time. Since you sent it, I have been working on finding reasons and applications for classic systems. Filling your slot in the disk cabinet would be a great place for an AMPRO or YASBEC systems. You might consider using it for a print spooler as the LATEX output looks like it might require lots of printer function commands.

However I might also recommend the Motorola trainers and sample boards. They were selling the 68HC16 demo board for $168. Not only do you get a full system, but they give you TONS of literature about the chip and support software (includes PCDOS based assembler and debuggers).

Being an Atari ST owner I know of their great cost/use ratio. I have MINIX for the ST and have found it lacking in some areas. Your MiNT sounds like something I would like to try, or better yet have you tell us all about it. I guess my main question is source code, does it come with it all?

You are also correct in that TCJ readers are interested in more than CP/M. I am finding out that some of our information is not available in any other magazine. That is one reason I keep saying we are and will continue to be the only magazine supporting classic systems. I am sure you have found out that most magazine are treating the Atari ST as a classic system as well. So if you want to do that article about Forth systems on ST's, we are ready and waiting. Thanks for all the comments. BDK.

October 5, 1992

Dear Bill:

I picked up the *TCJ* survey from GENIE yesterday and will answer by mail.

The first issue of *TCJ* under your leadership came the other day. I confess to a fair degree of disappointment in it. Having subscribed to the Journal since about issue 37, I have a great appreciation for what it has contributed to my knowledge and skill as a programmer in Z80. (My contributions to the Z-System community include ZDB, ZDT, ZBIB, and IOPZXR). I also enjoy other articles (although I am not at all into FORTH). My disappointment has more to do with the quality of the writing and editing than with the content. Your "Next Ten Years" on page 5 of issue #56 abounds with typos, grammatical and punctuation errors, and awkward, confused, and incomplete sentences. Maybe this does not matter to engineer types -- which I am not -- but a technical journal ought to be clear and accurate. I have some background as a writer, having written and reviewed for the late and lamented *PROFILES* and *MicroCornucopia*.

Now to answer the questions in your survey:

1) Why do I read *TCJ*? My main interest is in Z-Systems, ZCPR3, Z80 stuff, although I generally read everything else. I never miss Jay Sage's column. I miss terribly Bridger Mitchell's great contributions. Al Hawley and Terry Hazen are electronic friends whose articles are always helpful.
2) What do I want from *TCJ*? See #1.
3) How much equipment do I have? Kaypro II'83, modified with TurboRom, 42 Meg harddrive, modem, printer, Macintosh. I also use MS-DOS stuff at work.
4) Would I still subscribe to *TCJ* at $50/yr? $100/yr? Price of subscription will have very little bearing on my decision to renew when the time comes. Content of interest to me and quality and clarity of writing will determine.
5) Is there another magazine's format you like better? No comment.
6) What do I want? Quality articles -- in some cases longer may be better to cover subject adequately. I like regular writers who have something to say. I lean more toward software than hardware topics. Advanced and beginner topics are both appreciated. I was once a woeful beginner.
7) What do I hate most about *TCJ*? I think I have already covered this question except for one thing. I really hate the blue cover.

Bill, I know I've been fairly critical. I do want *TCJ* to succeed, and I'd be happy to help out in some way as a copy editor or consultant if that fits into your needs.

All the best to you and *TCJ*.

Joseph I. Mortensen, Midland, MI.

Thanks for your sincere comments Joseph.

It seems my panic in putting #56 together caused more errors than I thought. I am trying to cut them down, but must confess it will take several issues and freeing up of time to correct the problems. I have asked my contributing editors to act as reviewers of articles, but alas they too are very busy. In the near future I may call on you to review articles, or better yet help in getting them

and making sure they fully explain the topic for all levels of our readers.

We will be raising our rates after the first of the year. Mostly to offset and correct foreign mailing rates, but also to try and get some extra money so our writer might get some compensation for their hard work. As to your favorite writers, I only hope they see this and contact you to learn what you might like to see them comment and report on. I have asked all our past and present writers for articles, but it is up to them and not me as to when I have something of their's to publish.

You might consider writing something about your Z-System programs and using Z-products. Many of our readers are real beginners to CP/M and ZCPR and I am sure would love to hear about your early beginner projects. As always, I am grateful you took the time to respond to my survey. Thanks for your response and support of the Z-community. BDK.

Sept. 24, 1992

Dear Bill,

I have recently introduced a few people to Z-system over here, at the same time I have sent details of *TCJ* to them. Some of them have expressed an interest in *TCJ* but do not want to subscribe without seeing what it is like for themselves, they to far away from me to see one of my copies. I am enclosing a list of Z-system users and interested people, you might like to send them an introductory copy.

I have all *TCJ* from number 25 up to date and have recently ordered all the available back issues.

I own a SB180FX with ETS board and unused GT180 graphics board which I was hoping to use as a terminal but it is beyond me to get it sorted, I have the denied Turbo Module 2 with graphics toolbox, but it needs some software to get the lot together. Has anybody used this board? I have had no mention of it since it was introduced in Byte several years ago. I also have an Amstrad PCW8512 which runs Z3plus and doubles as a terminal for the SB180FX.

Keep up the good work with *TCJ* as without it I would have given up computers long ago or worse still use and IBM clone.

Your Sincerely, Mark Minting, Suffolk, England.

Thanks for the names Mark.

I sent all the people on your list a trial subscription. They should have gotten issue #57 by now, but knowing that the surface mail takes forever to get to Europe they could still be in shipment. I have decided to reprint (photo copy) back issues as sets, so you could get bound sets and fill in the gaps (if you have any).

It seems that CP/M is alive and well in England. I recently received a letter from a person still making and selling CP/M systems. The Amstrad appears to be very popular as well (something like 500,000 strong). As to your graphics board I can only hope someone reads this and send us the articles you need. The idea of using classic systems for terminals is not new, but certainly a great use. Maybe I need to find some articles about turning Kaypros into X-Terminals or ANSI compatible (MSDOS ANSI.SYS) terminals. With all the increase interest in UNIX, classic systems as terminals might be very cost effective.

Thanks again for the list of possible users, and keep up the good work in spreading TCJ's name around Europe. BDK.

Sept. 23, 1992

Dear Bill

I am enclosing my check for $20.00. This is in part a donation to support TCJ and for issue 56 which I have not received. Please send me a copy of issue #56. I don't understand the post office not forwarding TCJ since they forwarded everything else. Any way I don't want to miss an issue.

I have been an avid reader of TCJ since I met Art Carlson at the SOG in Bend many years ago. I started subscribing with issue 16, and have perused every issue since. The mix of articles has been great. I like both hardware and software subjects. Al Hawley's series on assembly language have been great. Jay Sage's have been near the top of my list.

What I have liked especially about TCJ is that in addition to articles I can understand there are articles that stimulate my though and interest. TCJ is essentially the only serious journal that is keeping alive CP/M, and 8 bit computing. Your idea of responding to the novice as well as advanced hobbyist is great. Please keep up the good work.

Yours sincerely, A. A. Straumfjord, Camp Sherman, Oregon.

P.S. I personally would find TCJ worth $50 or $100/yr but I fear many subscribers could not afford that.

Thanks for the extra money Al.

The last few months have been a big money drain for me personally. TCJ hasn't paid it's own way for a long time. I'll put some of your donation toward the increase in rates coming up next year. I have figured out it cost me about $2.75 to get an issue to you. So our problems with the post office was just too much. I tried sending make up issues out, but discovered I had spent almost an extra $100 doing that. I have had a talk with the post office and am sending TCJ with a minor change which should eliminate (or at least cut down) on incorrectly handled issues.

Well if you met Art at a SOG, you most likely saw me as well. I went to all the SOGs but the first one. I also helped Art man a table more than once. I really miss having SOG like meetings, so we are starting to consider having one in Sacramento next summer. Charles Stafford our Kaypro person has had several here already and wants to try and go national. With the Trenton computer

feast in New Jersey, only seems natural to have one on the west coast to match. I believe that 30% of my subscribers are from California. It may be hard getting your favorite writers out here, but we will try. I know they like seeing and hearing from their fans.

Thanks again, from Bill Kibler.

Oct. 3, 1992

Dear Bill,

I apologize for taking so long to reply to you. I,m terrible about writing letters. Articles I can whip up in a hurry, but letters...

I'm quite pleased with your new format. You probably share the kind of frustration I have felt with ASICs, since you specifically say none of those in articles. Thank you.

I have a little difficulty writing elementary level articles. The biggest problem is that no steering is possible. I often do classes in beginning LANS and in those I can adapt as I go.

There is really little difference between a stat mux and a LAN. That's usually the approach I take in trying to explain how a LAN works. I prefer not to confuse the issues of media versus access method. In effect, a multi-user computer becomes a degenerate LAN or in some cases a "LAN in a can".

As far as Netbios, leave me out. Most of the bad experiences I alluded to in my articles have been from PC networks. Personally I never thought any PC ever belonged on a LAN, although the BSD for PC might finally make it. Usually I tell people to get a workstation or some other capable machine and put two interfaces in it so that all the PC's can be isolated on their own segment without causing grief for everybody.

As far as network checkers, we make use of Netwatch quite a bit. It's freely available. There are some packages that are even better, but some of them have so much information that you can't read the display from a distance. Netwatch just scrolls the header of possible packets on the screen.

By the way, do you have Internet access? It was very convenient for me to be able to email my stuff to Jay Sage.

Wayne Sung

I loved getting your hand written letter Wayne. I am a lot like you when it comes to letters and articles. That is one reason I often hand write responses (only use computer for articles not letters), as well as being so much faster and personal. Please don't feel you can't "steer" our readers like you do your classes. I know what you mean about getting students to look over material and concepts by just prodding them in the right direction. However I find the students "steering" me more often than not. It is all those funny questions that make you think they just woke up that minute in your class even though you have been looking them straight in the face during the last hour of lecture. My advice to all my writers has been to think you are teaching a class and talking to students not readers. Maybe just tape recording your class and reprinting the questions and answers would make some great articles.

Novell's Netbios has been nothing but problems again. We thought we had the problem fixed only to find it back again last month. Yes my opinion about networks is changing but I am not burned out on them yet (soon however I am sure). I think you are a little overworked like most of the writers right now. The economy has everybody doing double duty and some like me doing triple duty.

I have looked for Netwatch but yet to find it. I will keep looking. How about an article on what and why of checking network data by using Netwatch? I know my students can't wait to learn about trouble shooting LANs. On ASICs, it really is PALs, they just burn out too often for my taste. Then when they do burn out, often the whole system becomes junk when the supplier is gone and you have no replacements for the part. The official stand at TCJ is do it first in TTL chips, then show the difference or alternate PAL option. Or in short give us choices and alternatives.

Lastly I am trying to get on Internet. I have joined Compuserve to access their Internet link, and now it seems GENIE has started doing Internet as well (which I already belong to.) With so many things to get up to speed on, Internet is really a back burner option at the moment. I hope to be using it some time after the first of the year. Thanks again for all you have done for TCJ. Bill Kibler.

10-17-92

Dear Mr. Kibler,

I've just received issue #57 of TCJ, with my article "Shell Sort in Forth". My only lament is that a shaky Undo Key apparently introduced a series of characters into a line of code in screen #7. The line between BEGIN and WHILE should be

DUP GAP @ - DUP 0< NOT SWAP S@ SV @ > AND

as in screen #13. Ah well, fast computers only make for faster mistakes!

Yours truly, Walter J. Rottenkolber

Thanks for the correction Walter. Your article fit in perfectly with the introduction to Forth. Not only was I able to give our readers a good how to get started, but also a great application in the same issue. I try to get supporting articles as much as possible, but that option is based more on what our writers give me and less on what I would like to see.

Our readers will be looking forward to one of your next articles. You have mentioned doing some more applications (game of LIFE, and Interrupt handling) which I am sure will go over very well. Your idea however of showing how to debug Forth code, especially after the typo above, is proably something our readers could use now. Learning to debug what you have written is always such a big hurdle for beginners to master. Thanks for all your support of TCJ. Bill Kibler.

The Z-System Corner

By Jay Sage

Regular Feature

ZCPR Support

Language Independence

In my last column I described a new idea for allowing message text in programs to be changed easily for adaption to various languages, or even just to suit a user's preferences. I would like to continue that theme this time as well.

First, in the process of working on the new ZFILER version 1.1, which includes the language overlay implementation, I thought of a couple of additional points. Second, Al Hawley sent me some very interesting comments, and he built on my ideas for use in the new version of BYE that he is writing. [BYE is the resident program that is used to implement what is called a "remote access system", a computer system that can be operated by a remote user via modem or direct connection to a serial port.]

An Addition to the Module Header

In the version of the text message (i.e., Z3TXT) module I finally arrived at by the end of the last column, there were several items included in a header at the beginning of the module. First, there was an opcode of "RST 0", which, if executed, would result in a warm boot. A file containing a Z3TXT module should never be executed, but at a cost of one byte we could protect ourself against that outside chance.

The header also contained the string of characters "Z3TXT" followed by a null (0) byte. Many Z-System modules include such identifiers. In this category are resident command packages (RCPs), flow command packages (FCPs), and environment descriptor modules (Z3ENVs). Programs, such as Bridger Mitchell's excellent JETLDR.COM, that load these modules from files into memory can use the ID string to validate the file, that is, to make sure that it is the kind of module that the user has stated it to be. User mistakes and damaged files can thus be detected. Some modules, however, such as named directory registers (NDRs) and terminal capability descriptors (Z3Ts), do not have such ID strings; they probably should have.

The next header item I included was the name of the program for which the Z3TXT file contains message text. Loading a text overlay intended for one program into a different program would probably not produce very satisfactory results <grin>. By including the program name in the header, a loader utility -- the equivalent of JETLDR -- would be able to catch such an error. Even if it didn't outright refuse to load the module, it could at least point out the possible problem to the user and ask for confirmation of the command.

The last item in the header I proposed was a three-letter language identifier. This would ensure proper identification of the language contained in an overlay module. This could help a user who could not otherwise identify the language with certainly, and it might also be of use to a loader utility.

There was one very important item that I omitted: the size of the module. If the loader cannot determine how much space has been allocated for the text module, it might end up installing a module that was too large. This would probably result in some executable code being overwritten. Again, the user would probably be quite disappointed by the results! This happened to me while I was working on ZFILER. I don't know if Al Hawley suffered any such accidents or whether he was just smart, but his list of suggestions to me included such an item for the header. The header, thus, now stands as follows:

```
        rst     0
        db      'Z3TXT',0 ; null-terminated ID
;               12345678 ; must be 8 characters,
        db      'PROGNAME' ; pad with spaces
;               123       ; must be 3 characters
        db      'ENG'     ; name of language
        dw      LENGTH ; length of module
```

The use of this length byte is somewhat tricky. When the program code is assembled, the messages for some language have to be included directly in the source code. In that case, the length word should contain the amount of space allocated in the program for the text overlay, not the amount of space actually used. It would probably be defined using an EQU directive in the source code.

On the other hand, when modules are assembled as independent overlays, the length word should contain the actual length of the module. Then, when the loader utility attempts to load a language overlay into a program's COM file, it could compare the length word in the overlay to the length word in the COM file to make sure that the new overlay will fit properly.

Here is the trickiest part. The loader must remember the original value of the length from the COM file and reinstall that value after the new language overlay has been installed. In other words, the length word in the header serves two distinct purposes: in a free-standing overlay file it represents the space needed for the module; once installed into a COM file it represents the space available for the module.

New Data in the Text Module

In my examples last time, the Z3TXT module contained only null-terminated text strings. While working on the text module for ZFILER, I realized that one might wish to include some other types of data. Al Hawley also showed me an additional way to handle the address table and data. I will now describe both of these extensions.

In ZFILER, as in many other programs, users are occasionally prompted with a question that requires a yes or no answer. In the past, the answer has been given by pressing either the "Y" key or the "N" key. This is all well and good for English speakers, but other languages may use different words represented by different letters. (Interestingly enough, though the word for "no" varies, the first letter is nearly universally "N" for European languages; "Y", on the other hand, will often not fit the bill.)

I wanted to fix this in ZFILER, so I decided to put the two letters representing an affirmative and a negative answer into the text overlay and to have the code compare the user's response to the variable characters stored there. Recall the complete structure of the module as I proposed it last time. Between the header and the actual message strings was a table of address offsets to the text strings. The appearance was thus as follows:

```
z3txt:                  ; start of module
; header goes here (see above)
        ...
; table of address offsets
msg1:   dw      .msg1 - z3txt
msg2:   dw      .msg2 - z3txt
        ...
; actual message strings
        ...
```

Since the length of the "yes" and "no" characters is fixed, namely one character, there is no need to use indirect addressing as with the variable-length message strings. The letters can be put directly into the address-offset table. It might then have the following appearance:

```
; table of address offsets and
; ..fixed-length data
affirm:   db      'Y'
negate:   db      'N'
msg1:     dw      .msg1 - z3txt
msg2:     dw      .msg2 - z3txt
```

In Al Hawley's new version of BYE for Z-Systems (called NZBYE) there are places where one wants the code to modify a string dynamically. For example, there might be a message of the form, "You have 7 minutes remaining." This message might be displayed each minute once the user has fewer than 10 minutes of allowed access remaining for the current call. Here is how Al sets up the module.

```
; table of address offsets
timmsg: dw      .timmsg - z3txt
timval: dw      .timval - z3txt
msg2:   dw      .msg2 - z3txt
        ...
; actual message strings
  db    ' minutes remaining',cr,lf,0
        ...
```

The table values stored at the addresses labeled "timmsg" and "msg2" point to the beginnings of complete message strings (at addresses ".timmsg" and ".msg2"). Those messages would be displayed as I described last time using a special string-printing subroutine. In Al's extension, the address at label "timval" would be used differently. A fixed-length character string would be computed by the program and written into the space at label ".timval", whose address would be computed from the table entry at "timval". A slightly more elaborate version of such a subroutine is presented later in this column.

Alexander Schmid and I faced a similar problem in ZFILER but handled it differently. We broke the strings into separate pieces. Sticking with the above example, we would have:

```
; table of address offsets
timmsg:   dw      .timmsg - z3txt
timmsg2:  dw      .timmsg2 - z3txt
msg2:     dw      .msg2 - z3txt
          ...
; actual message strings
          ...
```

In the program code, the two pieces of the message would be printed separately, each by a call to the message-printing subroutine. The time value would be sent to the terminal in between those two calls. It is not clear to me which method is best. The second approach has an extra null byte at the end of the first message string, and the display subroutine has to be called an extra time. On the other hand, with Al's method, the address has to be calculated and the computed value has to be stored into the message string instead of just sending it directly to the screen. My feeling is that the two methods are so nearly equal in efficiency that the choice should be made based on one's programming style preference.

Some Additional Points

Al Hawley was still not completely satisfied with the storage efficiently in the text module. Remember that last time we began with an approach in which each individual message started at a fixed address, with extra space included with each one to allow for longer messages in another language. Then we added the address-offset table -- with all entries in fixed locations -- so that the messages could be of different lengths without having to waste expansion space for each one.

With this approach, however, one still has to allocate some overall extra space in case the complete collection of messages in another language is longer. Consequently, space is wasted in the COM file when a shorter language overlay is loaded. This bothered Al.

I had thought about this, too, and had concluded that there really was nothing one could reasonably do about it. If the Z3TXT module could be placed at the end of the COM file, then its length could be variable. However, programs generally store data after the end of the material stored in the COM file. If the Z3TXT module were appended, then the code would have to provide indirect addressing for all data references, with the addresses depending on the length of the currently loaded Z3TXT module.

This approach would also cause prob-

lems with type-4 programs. These are programs whose load-address is determined at run-time by the ZCPR34 command processor. Generally, type-4 programs are loaded as high as possible in memory. This requires knowing how much memory the program uses. Variable-length text overlays stuck on the end of the COM file would greatly complicate this situation.

Al Hawley's thinking on the subject was colored by the task he was tackling. BYE is not a transient (temporary) program. The COM file (BYE.COM) causes a block of code called a resident system extension (RSX) to be installed semi-permanently at the top of memory under the normal CP/M operating system code. There it intercepts operating system calls, changing existing functions and adding new ones. Because it is resident, BYE reduces the memory space available to all other programs. Therefore, it is of great importance to keep the code as short as possible. To accomplish this with NZBYE, Al came up with a very clever idea that I will describe in just a moment.

First I want to mention one other new consideration that arose in connection with NZBYE. Remember that BYE's function is to make the modem port an extension of the system's console (keyboard and screen) so that the computer can be operated either by the local operator (the sysop) or a remote caller.

Picture Al sitting at the console of his Ladera Z-Node in Los Angeles while a caller from France is on-line. Being an accommodating person, Al would like user messages sent out by BYE to appear in French. On the other hand, there are some messages generated by BYE that are for the sysop only and appear only on the local screen. These should be in English for Al. For Helmut Jungkunz in Germany, however, they should be in German, even when the French caller is using the system.

Well, the conclusion Al came to was that NZBYE needed to break the messages into two groups: local messages and remote messages. (The remote messages go to the local console as well. Ideally, these messages would be provided in two languages and separate versions would be sent to the local and remote consoles. However, this would waste memory by complicating the code and increasing the space taken up by the message text.)

Since each user of BYE has to assemble the program anyway (because of the choice of various options and system characteristics), there is no reason not to assemble in the message text as well. To make it easy to change the messages, Al has gathered them into one place and put their code in a separate LIB file. Those messages do not have to be addressed via an address-offset table; they can be addressed directly. Conditional assembly pseudo-ops can be included so that when certain optional functions are omitted from the BYE code, the corresponding messages can be omitted as well.

The remote messages have to be loaded dynamically depending on the caller. They could be handled as we have described for standard programs, but then one would have to pre-allocate enough space to accommodate the longest set of messages for all supported languages. Al conceived the following approach.

The remote message text is not stored in BYE at all; it is stored in a separate RSX that is loaded after (underneath) BYE. As a result, its size can vary. Here is a more detailed description of what happens. When BYE.COM is run, it has to deal with two tasks. As before, it has to make sure that the BYE RSX is loaded, installing it in upper memory if it is not already present. Then it has to make sure that the correct language RSX is loaded. If BYE.COM has just loaded the BYE RSX, then it has to load the required language module, too. If BYE was already present in memory, then BYE.COM has to remove the old language RSX and then install the one now required. (It could try to determine if the right one is already installed, but it may be easier to just dump the old one and start over.)

There is now one new complication. With our Z3TXT modules, the main code does not know where the actual message strings are stored, but it does know where the address-offset table is stored. This may not be the case with the BYE.RSX. We could change the structure of the language RSX, placing the address offset table at the end of the module. This would be a fixed address relative to the BYE RSX code, and the address could be adjusted just as all other addresses are adjusted when a relocatable module, such as an RSX, is loaded. However, one would like to maintain compatibility and minimize the proliferation of standards.

In order to use the same module structure, Al added an extra level of address indirection. Since BYE.COM handles the installation of both RSXs, it has the global picture and can provide the BYE RSX with the address at which the language RSX begins. This address is loaded into a specific data word in the BYE RSX. Here is what Al's double-indirection implementation looks like.

Suppose the language RSX is to contain the messages "msg1" and "msg2" that we showed in the earlier examples. An RSX contains some additional header code for intercepting operating system calls, for protecting itself (so it does not get removed during warm boots), for identifying itself, and for removing itself from memory. The remainder of the RSX will look just like the Z3TXT modules used for other programs. The starting address of that code ("z3txt") is what BYE.COM will store at, say, the label "byerem" in the BYE RSX code.

Al uses a subroutine to compute the real addresses of the messages through two levels of indirection. This sounds more complicated than it really is; I hope you will try to follow it. When the routine is called, the HL register pair has been set to the relative position (or offset) in the Z3TXT address-offset table for the desired message (or other data structure). For example, in the case of the first message, this value would be (msg1-z3txt). Although the individual values of "msg1" and "z3txt" will vary depending on where the language RSX is loaded in memory, the difference is a constant that is known at the time BYE is assembled. Since "msg1" is the first

message in the table, the value would be 20 (the length of the Z3TXT module header). The second message would have the value 22, and so on.

Here is the complete code for Al's subroutine, called GETM2. The value of the offset to the offset in HL is converted into the absolute address of the specified object and returned in HL. Otherwise, only register A is altered; if one desired, register pair AF could be protected on the stack as is done with DE.

```
  push    de          ; save DE on stac
  ld      de,(byerem) ; address of z3txt
; first indirection: get absolute address where
; offset to message text is stored and load the
; value into HL.
  add     hl,de       ; real address of offset
  ld      a,(hl)      ; get low byte into A
  inc     hl          ; point to high byte
  ld      h,(hl)      ; get high byte into H
  ld      l,a         ; get low byte into L
; second indirection: given offset in HL, calculate
; actual address of the message text in Z3TXT
module.
  add     hl,de
  pop     de          ; restore DE
```

Plans for Future Columns

In two issues, TCJ will be marking its 10th anniversary. At the request of the editor, for that issue I will be revisiting the topic of getting Z-System running on a computer that is currently running a form of standard CP/M. The column will be addressed to novices, but there might be some tricks and thoughts of interest even to veteran Z-System users.

To learn about the subject myself, I plan to dig out one of the many Kaypros warehoused in my basement and go through the complete process of bringing up an NZCOM Z-System on it. I will also install Z3PLUS on one of my CP/M-Plus computers, perhaps the Televideo 803, on which I currently have CP/M-2.2 and NZCOM but for which I believe I have a CP/M-Plus boot disk as well.

For next time I hope to describe some advanced uses of the PMATE/ZMATE text editor. MATE is generally used manually to edit files, but it can also be used as a generalized, programmable text-processing tool. I have installed a special permanent macro that greatly facilitates this kind of use.

For example, on my 486/33 PC at work I have been running many long series of automated simulations of electronic circuits using PSPICE. A 4DOS batch file starts the process by writing out a file with the circuit parameters (for example, capacitor values, voltage levels, clock rise and fall times). It then invokes PSPICE and feeds the output file to PMATE, which uses a macro to analyze the data and determine whether or not the circuit operated as desired. This macro is stored in a separate file (or files, if there are subroutine macros). The results of the evaluation are written out in files that are then read by the 4DOS batch file and used to generate the next set of simulation parameters. In this way, I can leave the computer running unattended for days or even weeks, and when I come back there is a nice report awaiting me with a summary of the results. Readers may be able to think of many uses for this approach.

Language Independence is a big problem. While checking on new versions of Novell's Netlite, I discoverd they have a separate version for each language. It solved the problem but at a price in programming time and stocking different versions. Jay's ideas sure make more sense and like Al Hawley's concept of needing different languages for different sections. From a maintenance stand point, by having different modules you could load them on the fly, that would allow english speaking support people to work at a site that normally uses another language.

I am sure there are many other options to this problem. If you have one, or found some "slick" way of enhancing Jay's work, let Jay know.

Contact Jay on GENIE as JAY.SAGE or see his add on the inside cover for other options. BDK.

In Issue #59

- Programming the 6526 Communications Adapter. This in-depth article shows how to approach and program the adapter in BASIC. A great review and startup articles for those wanting to learn the insides of a computer.

- D/A Conversion on the Cheap. This is part two of a series on how to generate analog signals using a few resisters and an op-amp. Programming support using Forth and any parallel port makes this one of our platform independent projects.

- Real Computing tackles losing the "Superblock" on a MINIX disk system.

- Z-System covers PMATE, one of Jay's favorite programs.

- More articles on S-100 and Kaypros as well.

Support Requests

To request support or assistance on a given project, please write our support people directly. TCJ's Technical Editors place their address in their columns so that you can get faster response by writing to them directly. Your requests will be included in later articles. Our editors reserve the right to refuse service or redirect you to a more appropriate source. Please include as much background information as possible.

Contacting TCJ
Mail:
TCJ
P.O. Box 535
Lincoln, CA 95648

On GENIE:
B.Kibler

Real Computing

By Rick Rodman

32-Bit Systems

All Readers

MINIX, UZI and GNU

Minix, Uzi, Gnu - and lessons from home movies

Minix goings-on

Many Minix OS hobbyists have changed over to Linux, for two reasons: first, there are no restrictions to worry about, and second, Linus Torvalds, the author, is in favor of it becoming a big, complex operating system.

The Minix community has been revitalized by a sudden new burst of activity, however. A gentleman named Frans Meulenbroeks proposed that the Minix kernel be rewritten as a clean microkernel design. As Andy Tanenbaum pointed out with respect to Linux, the microkernel type of OS design has been almost universally acclaimed as superior, and most new operating systems use that design. While Minix started out as a microkernel, the severe design constraints of the Intel 8086 architecture compromised it to the point that most advantages of a microkernel were lost. The new design discussion has basically centered around building a true microkernel design while still retaining as much as possible of Minix.

In a true microkernel, the protected kernel of the operating system does task switching and message passing *only*. Device drivers and other service tasks run as user-level tasks.

Frans introduced his proposals in an essay entitled "Proposal to restructure MINIX". Here are a few excerpts.

"Minix as it is now is quite cumbersome if you want to add new tasks. You have to change NR_TASKS, table.c, add the new task, recompile fs/mm/kernel/tools (they all depend on the include file which defines NR_TASKS). Quite a job.

"This way it is quite complicated to add/remove drivers. Furthermore it is difficult to understand how the system works.

"My idea is to restructure the kernel into a number of different processes. Each process corresponds with a kernel task (e.g. the floppy driver). None of them accesses variables from other task (low coupling).

"To support this, there is a microkernel (further on called core, to avoid confusion with the current kernel) mainly consisting of system.c and proc.c (and perhaps a part of main.c).

"This core should have the following functions:
- context switching
- message passing
- basic clock handling
- memory copying between tasks
- support for interrupt handling
- creation and removal of processes.

"Functions supplied by the core itself are:
- the functions from system.c and proc.c
- install interrupt handler
- deinstall interrupt handler
- attach to major device number (this way device numbers can be connected to drivers). This can simply be the mapping of the device number to the task id.

"If we identify tasks by a bit in their proc struct and make them standalone, it is very easy to add a new task to the system or remove one. Of course only the super user can do so (by means of a new system call). An implementation of this could be done stepwise by:
- introducing task and driver bits in the proc struct. perhaps also a bit must be reserved to identify the idle task.
- rewrite the macros in proc.h to use these driver bits instead of the NR_TASKS.
- rewrite the code which uses things like END_TASK_ADDR.
- allow message passing between all drivers.
- take a simple driver and rewrite it so that it is stand alone compilable (so it uses the sys_copy system call instead of directly calling phys_copy and umap). The printer or the memory task seems a good candidate for this.
- add a way to attach/detach an interrupt routine to an interrupt source
- add a way to attach/detach to a major device number
- add a system call which lets a program start as task or server.
- remove the rewritten driver from the system and start it using the new system call
- add code to the build process to load standalone drivers on boot time
- rework the other drivers, taking them out of the kernel one at a time
- clean up the remaining kernel (which should mainly be system.c proc.c and some hulp [sic] code)."

There have been many related discussions on Minix' message system, which is quite cumbersome due to the six message formats. Some have suggested using variable-length messages. Of interest is the intent to keep Minix' size small and manageable - in fact, to make its kernel even smaller and simpler than it is today. Even Andy Tanenbaum, the author of Minix, contributed to the dis-

cussion in a positive way. While Minix 1.6 will be released soon as an incremental improvement over 1.5, it's good to see that the community has rediscovered a technical vision for Minix' future - a simple but extensible operating system that can be completely understood and extended in meaningful ways by a single individual.

Recently published on Usenet was a syntax guide for the PC MINIX assembler. This assembler doesn't follow the arcane and stultified "structure" conventions of Intel's and Microsoft's assemblers. Nor does it follow the more normal but still quirky conventions of the CP/M-86 assembler. Instead, it follows those of IBM's PC/IX assembler. "Aha! Of course!" I hear you all saying. Well, for those of you who have been on the other side of the planet from such earth-shaking products as PC/IX (like me) and don't know that syntax like the back of your hand, send me a quick note and I'll send you a copy. It's too lengthy to print in this column.

Uzi

Another free operating system is Uzi, developed by Doug Braun for the Z-80 processor. Written in C, it is a fairly complete Unix Version 7 clone - in a little over 6,000 lines of code. Although it's a monolithic kernel design, it's pretty small. It takes 32K bytes of RAM for the kernel itself on the Z-80. I've been thinking about porting it to the NS32 processor - and putting it in PROM.

Gnu tools under OS/2

GCC, the Gnu C Compiler, and G-Plus, the Gnu C++ translator, have been ported to OS/2 2.0. Some folks have been using them for code development. This is a rocky road to travel down, but the product is free (at least if you have ftp access). Interestingly, the 32-bit linker (LINK386) and the Resource Compiler (RC) are supplied with the operating system, so all you need is the compiler itself. I don't have copies of this software yet, but I'm looking to acquire them shortly. Other Gnu tools, such as Gawk (Gnu AWK), Bison, and Gnu Emacs are being ported and will be available shortly.

I mentioned last time that IBM presently offers only C for OS/2 2.0. There are other compilers which can generate 32-bit code for the system, notably Watcom C, Microway's Fortran, and Logitech Modula-2. IBM also provides a strange programming environment which works with C, called System Object Model (SOM). It's heavily into object-oriented mumbo-jumbo. In response to much hue and cry, IBM has promised to release a C++ compiler "soon". Personally, I don't care. C lets me do what I need to do.

Programming is the art of effectively expressing algorithms in a form comprehensible to both machines and humans. That language which does so most efficiently for a specific class of problems is the best for that class. In my investigations so far, object-oriented programming languages I've looked at have failed to either generate efficient code or express algorithms clearly. The industry will probably never make a transition to OOP, but OOP may be more than a fad nonetheless; it may be an incremental step to the next technology.

Remember home movies?

You know, Super 8 and Standard 8? They were big once. They had full color, slow-motion, sound, and zoom. What happened to it all? It disappeared overnight. Video equipment blew it away - even faster than the video proponents themselves expected.

OK, in retrospect, we can say that the technology was lousy. You only got about 3-1/2 minutes per roll of film. You had to wait for it to be developed. Especially with Standard-8, you could mess up and lose the whole roll trying to unload the film.

The lesson we learn from the rapid demise of home movies is that people will dump a technology almost instantly when something better comes along - especially one that addresses their long-ignored frustrations and problems.

In TCJ #57, I discussed some technology fads I expect real soon. Some folks have pointed out that I forgot to mention still video. Still photography has terrible technology problems - the same ones as home movies, in fact. I've mentioned losing whole rolls of film due to camera jams, for example, to other photographers, and been surprised to hear that even professional photographers have these problems. Besides, the insurance industry is adopting still video like mad. With a video digitizer boards, the photos can be centrally stored on computers and made available via LANs.

Based on the lessons learned from the home movie to VCR transition, some have predicted that still video will completely eradicate film even while the cost is higher.

Personally, I don't think so. The cost is still way too high for the full suite of necessary equipment - don't forget, you've got to include printing hardware. Also, the resolution still isn't too good. The day when Hasselblads sit unsold on flea-market tables at giveaway prices is still a couple of years off.

Nevertheless, this may be the best time for technical folks to familiarize themselves with the technology. With industry and government moving away from film so rapidly, there will be lots of opportunities in file conversion, viewing/cataloging software, and suchlike. One nice thing coming out of the imaging technology drive is the move to SCSI as a standard I/O interface. Most high-end color printers, for example, have SCSI interfaces, as do most of the new scanners coming out.

Next time

Push the edge of the envelope with us as we try to fix corrupted Minix filesystems, using only our bare hands, the ROM monitor, and RDMINIX tools under DOS. Where else but in computers

would something so fragile be called a "superblock"?

Where to call or write:

BBS or Fax: 1-703-330-9049 (There's an autoswitch feature in the fax machine; strange though this seems, it appears to work just fine.)

Here is a brief synopis of the MINIX assembly language. It is the same as IBM's PC/IX assembler.

2. TOKENS
2.1 Numbers
Same as C

2.2 Character Constants
Same as C, supporting \n, \t, \b, \r & \f

2.3 Strings
Same as C

2.4 Symbols
May contain any letter, digit, ".", "~" or "_", but cannot have a digit or "~" as the first character.

All global names have 8 significant characters

The names of the 8086 registers are reserved ([abcd][xlh], [cdes]s, [ds]i, [sp]p, bx_[ds]i & bp_[ds]i). The last two forms indicate register pairs; these names are used in the "base + index" addressing mode (section 6.1).

Names of instructions and pseudo-ops are not reserved. Alphabetic characters in opcodes and pseudo-ops must be in lower case.

2.5 Separators
Commas, blanks, and tabs are separators and can be interspersed freely between tokens, but not within tokens (except string and character constants) or between the tokens of an expression. Commas are only legal between operands.

2.6 Comments
The command character is "|".

2.7 Opcodes
Listed below.
Notes:
1) Different names for the same instruction are separated by "/".
2) Brackets ([]) indicate that 0 or 1 of the enclosed characters can be included.
3) Curly braces ({}) work similarly, except that one of the enclosed characters must be included.

2.7.1 Data Transfer

2.7.1.1 General Purpose

mov[b]	dest, source	| Move word/byte
mov{bw}	dest, source	| Move word/byte from source to dest
pop	dest	| Pop stack
push	source	| Push stack
xchg	op1, op2	| Exchange word/byte
xlat		| Translate

2.7.1.2 Input/Output

in[w]	source	| Input from source I/O port
in[w]		| Input from DX I/O port
out[w]	dest	| Output to dest I/O port
out[w]		| Output to DX I/O port

2.7.1.3 Address Object

lds	reg,source	| Load reg and DS from source
les	reg,source	| Load reg and ES from source
lea	reg,source	| Load effect address of source to reg | and DS
seg	reg	| Specify seg reigster for next | instruction

2.7.1.4 Flag Transfer

lahf		| Load AH from flag register
popf		| Pop flags
pushf		| Push flags
sahf		| Store AH in flag register

2.7.2 Arithmetic
2.7.2.1 Addition

aaa		| Adjust result of BCD addition
add[b]	dest,source	| Add
adc[b]	dest,source	| Add with carry
daa		| Decimal Adjust acc after addition
inc[b]	dest	| Increment by 1

2.7.2.2 Subtraction

aas		| Adjust result of BCD subtraction
sub[b]	dest,source	| Subtract
sbb[b]	dest,source	| Subtract with borrow from dest
das		| Decimal adjust after subtraction
dec[b]	dest	| Decrement by one
neg[b]	dest	| Negate
cmp[b]	dest,source	| Compare
cmp{bw}	dest,source	| Compare

2.7.2.3 Multiplication

aam		| Adjust result of BCD multiply
imul[b]	source	| Signed multiply
mul[b]	source	| Unsigned multiply

2.7.2.4 Division

aad		| Adjust AX for BCD divison
cbw		| Sign extend AL into AH
cwb		| Sign extend AX into DX
idiv[b]	source	| Signed divide
div[b]	source	| Unsigned divide

2.7.3 Bit Manipulation
2.7.3.1 Logical

and[b]	dest,source	| Logical and
not[b]	dest	| Logical not
or[b]	dest,source	| Logical inclusive or
test[b]	dest,source	| Logical test
xor[b]	dest,source	| Logical exclusive or

2.7.3.2 Shift

sal[b]/shl[b]	dest,CL	| Shift logical left
sar[b]	dest,CL	| Shift arithmetic right
shr[b]	dest,CL	| Shift logical right

2.7.3.3 Rotate

rcl[b]	dest,CL	| Rotate left, with carry
rcr[b]	dest,CL	| Rotate right, with carry
rol[b]	dest,CL	| Rotate left
ror[b]	dest,CL	| Rotate right

2.7.4 String Manipulation
The following instructions address source strings through SI and dest string through DI.

cmp[b]	| Compare
cmp{bw}	| Compare
lod{bw}	| Load into AL or AX
mov[b]	| Move
mov{bw}	| Move
rep	| Repeat next instruction until CX=0
repe/repz	| Repeat next instruction until CX=0 and ZF=1
repne/repnz	| Repeat next instruction until CX!=0 and ZF=0
sca{bw}	| Compare string element ds:di with AL/AX
sto{bw}	| Store AL/AX in ds:di

2.7.5 Control Transfer
Displacement is indicated by opcode; "jmp" generates a 16-bit displacement, and "j" generates 8 bits only. The provision for "far" labels is described below.

As accepts a number of special branch opcodes, all of which begin with "b". These are meant to overcome the range limitations of the conditional branches, which can only reach to targets within -126 to +129 bytes of the branch ("near" labels). The special "b" instructions allow the target to be anywhere in the 64K-byte address space. If the target is close enough, a simple conditional branch is used. Otherwise, the assmebler automatically changes the instruction into a conditional branch around a "jmp".

The English translation of the opcodes should be obvious, with the possible exception of the unsigned operations, where "lo" means "lower", "hi" means "higher", and "s" means "or same".

The "call", "jmp", and "ret" instructions can be either intrsegment or intersegment. The intersegment versions are indicated with the suffix "i".

2.7.5.1 Unconditional

br	dest	| jump, 16-bit displacement, to dest
j	dest	| jump, 8-bit displacement, to dest
call[i]	dest	| call procedure
jmp[i]	dest	| jump, 16-bit displacement, to dest
ret[i]		| return from procedure

2.7.5.2 Conditional with 16-bit Displacement

beq	| branch if equal
bge	| branch if greater or equal (signed)
bgt	| branch if greater (signed)
bho	| branch if higher (unsigned)
bhis	| branch if higher or same (unsigned)
ble	| branch if less or equal (signed)
blt	| branch if less (signed)
blo	| branch if lower (unsigned)
blos	| branch if lower or same (unsigned)
bne	| branch if not equal

2.7.5.3 Conditional with 8-bit Displacement

ja/jnbe	| if above/not below or equal (unsigned)
jae/jnb/jnc	| if above or equal/not below/not carry (unsigned)
jb/jnae/jc	| if not above nor equal/below/carry (unsigned)
jbe/jna	| if below or equal/not above (unsigned)
jg/jnle	| if greater/not less nor equal (signed)
jge/jnl	| if greater or equal/not less (signed)
jl/jnqe	| if less/not greater nor equal (signed)
jle/jgl	| if less or equal/not greater (signed)
je/jz	| if equal/zero
jne/jnz	| if not equal/not zero
jno	| if overflow not set
jo	| if overflow set
jnp/jpo	| if parity not set/parity odd
jp/jpe	| if parity set/parity even
jns	| if sign not set
js	| if sign set

2.7.5.4 Iteration Control

jcxz	dest	| jump if CX = 0
loop	dest	| Decrement CX and jump if CX != 0
loope/loopz	dest	| Decrement CX and jump if CX = 0 and ZF = 1
loopne/loopnz	dest	|Decrement CX and jump if CX!=0 and ZF=0

2.7.5.5 Interrupt

int	| Software interrupt
into	| Interrupt if overflow set
iret	| Return from interrupt

2.7.6 Processor Control
2.7.6.1 Flag Operations

clc	| Clear carry flag
cld	| Clear direction flag
cli	| Clear interrupt enable flag
cmc	| Complement carry flag
stc	| Set carry flag
std	| Set direction flag
sti	| Set interrupt enable flag

2.7.6.2 External Synchronisation

esc	source	| Put contents of source on data bus
hlt		| Halt until interrupt or reset
lock		| Lock bus during next instruction
wait		| Wait while TEST line not active

3.THE LOCATION COUNTER, SEGMENTS AND LABELS
3.1 Location Counter
The special symbol "." is the location counter and its value is the address of the first byte if the instruction in which the symbol appears and can be used in expressions.

3.2 Segments
There are three different segments: text, data and bss. The current segment is selected using the .text, .data or .bss pseudo-ops. Note that the "." symbol refers to the location in the current segment.

3.3 Labels
There are two types: name and numeric. Name labels consist of a name followed by a colon (:). Numeric labels consist of one or more digits followed by a dollar ("$"). Numeric labels are useful because their definition disappears as soon as a name label is encountered; thus numeric labels can be reused as temporary local labels.

4. STATEMENT SYNTAX
Each line consists of a single statement.

4.1 Null Statments
Contain neither an assembler command nor a pseudo-op. May contain a label or comment.

4.2 Instruction Statements

label:	opcode	operand1, operand2	| comment

4.3 Pseudo-op Statements
An assembler instruction, see below.

5. EXPRESSIONS
5.1 Expression Syntax
5.2 Expression Semantics
The following operator can be used: + - * / & ! < (shift left) > (shift right) - (unary minus) \ (unary complement). 32 bit integer arithmetic is used. Division produces a truncated quotient.

6. OPERAND SYNTAX
6.1 Addressing Modes

8-bit constant	mov	ax, *2
16-bit constant	mov	ax, #12345
direct access (16 bits)	mov	ax, counter
register	mov	ax, si
index	mov	ax, (si)

index + 8-bit disp.	mov	ax, *-6(bp)
index + 16-bit disp.	mov	ax, #400(bp)
base + index	mov	ax, (bp_si)
base + index + 8-bit disp.	mov	ax, *14(bp_si)
base + index + 16-bit disp.	mov	ax, #-1000(bp_si)

Any of the constants or symbols may be replacement by expressions. Direct access, 16-bit constants and displacements may be any type of expression. However, 8-bit constants and displacements must be absolute expressions. Floating point instructions that explicity reference the floating point register stack do so by specifying stack offsets:

```
fadd    0, 3        | Add fourth to top
fxch    2           | Exchange second and top elements
```

Arithmetic instructions must have two operands as in the example. Other instructions have only one. Floating point arithmetic instructions that mimic a stack machine can be coded with no operands. For example

```
fadd
faddp   1,0         as the same.
```

6.2 Call and Jmp

With the "call" and "jmp" instructions, the operand syntax shows whether the call or jump is direct or indirect; indirection is indicated with an "@" before the operand.

```
call    _routine    | Direct, intrasegment
call    @subloc     | Indirect, intrasegment
call    @6(bp)      | Indirect, intrasegment
call    (bx)        | Direct, intrasegment
call    @(bx)       | Indirect, intrasegment
calli   @subloc     | Indirect, intersegment
calli   cseg, offs  | Direct, intersegment
```

Note that call (bx) is considered direct, though the register isn't called, but rather the location whose address is in the register. With the indirect version, the register points to a location which contains the location of the routine being called.

7. PSEUDO-OPS

7.1 Assigment

Either using the symbol as a label when it is set to "." for the current segment with type relocatable. Or via symbol = expression when symbol is assigned the value and type of its arguments.

7.2 .long, .word and .byte

These commands take one or more operands, and for each generate a value whose size is a long (4 bytes), word (2 bytes) or a byte.

7.3 .ascii and .asciz

These commands take one string argument and generate the ASCII character codes for the letters in the string. asciz automatically terminates the string with a null (0) byte (a C string?).

7.4 .zerow

.zerow take one argument and generates that number of zero words

7.5 .even

.even generates a null byte is the location counter is odd, thus making it even.

7.6 .text, .data and .bss

These three commands select the current segment. The assembler always begins in the .text segment. No code can be assembled in the .bss segment.

7.7 .globl

.globl declares that each of its operands, which must be names, are globally visible across all files of the program. The names need not be defined in the current file, but if they are, their type and value arise from that definition independently of the .globl declaration.

7.8 .comm

.comm declares storage that can be common to more than one module. There are two arguments: a name and an absolute expression giving the size in bytes of the area named by the symbol. The type of the symbol becomes external. The statement can appear in any segment.

-*-*-*

Affordable Microprocessor Development Tools

by Tim McDonough

Embedded Systems

Beginners and Up

Programming Tools

Recently, I've received several letters from readers asking questions about past articles and questions in general about working with embedded systems. I really enjoy hearing from you because it lets me know someone out there is interested and it helps me figure out what direction to take for future articles. At the end of this article, I've listed several ways that you can contact me. My preference, if you have the capability, is via one of the two forms of electronic mail. In most cases I answer e-mail the day I receive it. Besides, it saves trees and helps keep paper from piling up around here.

The most frequently asked question I get is "Where can I get affordable development software for the (fill in the blank) microprocessor?" Affordable can mean different things to different people. In the case of the hobbyist or small business it typically means the cheapest thing that will get the job done.

There are a lot of companies that produce assemblers, compilers, and interpreters for various microprocessors and microcontrollers. I'm not going to attempt to list them all because I'm sure I'd miss someone with a great product. The tools I use are all MSDOS programs since the PC clone is my development platform. To my knowledge the products mentioned below are only available for that platform.

A good source of reliable, solid cross-assemblers, simulators, and disassemblers is PseudoCorp. The Level 1 demonstration versions of many of their cross-assemblers are available on many BBS systems and from major shareware vendors. I've personally used the commercial (Level 2) 8051, 68HC11, and 8096 products. Cost is about $50 each for the Professional Level 2 Assemblers.

The PseudoCorp products were among the first that I used. Those of you that have followed my 8031 articles here may recall that the program listings mention them. They offer features such as include files, a powerful flexible macro facility and the ability to generate segmented object files for systems where your code and/or data spans several ROMs. In writing code for the TCJ articles I avoided using these features to make it as easy as possible for you to use the software of your choosing.

Over the last few years there has been more and more interest in using high level languages, especially C, for embedded systems development. For the hobbyist or the small developer that meant drooling over compilers that cost anywhere from $800 to well over $2,000 in some cases. I don't know what your hobby budget is like but for most people a $2,000 compiler is probably out of the question.

A couple of years ago I ran across Dave Dunfield of Dunfield Development Systems on a BBS. Dave has a remarkable collection of microprocessor software development tools that seems to grow each month. His flagship products are MICRO-C, a very portable C compiler, and his XASM series of cross-assemblers. Currently his products are available for the 6800, 6801/6803, 6805, 6502, 6809, 68HC11, 8051/8052, 8080/8085/Z80, 8086 and 8096 microprocessors. Versions for the 68HC16 should be available by the time you read this.

So, just what does affordable mean? Dunfield Development Systems sells a package that includes ALL of the assemblers mentioned above with documentation on disk for $49.95 plus shipping. If you plan to delve into multiple processors and are going to stick with assembly language you probably won't find a better deal than Dave's XASM product.

Sooner or later your projects will grow in complexity and unless there's an overriding need for flat out speed you'll probably yearn for a high level language to use for your projects. If you like C then Dunfield's MICRO-C for the processor of your choice is probably the best bargain you're likely to find in the embedded systems market. The "Developer's Kit" for a particular processor has just about everything you need to write code except the microprocessor manufacturer's data book. My MICRO-C 8051 package included the following items: MICRO-C 8051 compiler, ASM51 cross-assembler, optimizer, hand coded (in assembly) standard library functions (with source code), monitor/debugger (with source code), macro pre-processor, linker and a text editor. The best news--a "Developer's Kit" for any of the supported processors is only $99.95 plus shipping when you buy it direct from Dunfield Development Systems.

MICRO-C is a K&R style C compiler. It provides features of special interest to the embedded systems developer such as being able to write interrupt handlers in C and extending the standard library with some additional bit manipulation functions. Other features include inline

assembly language capabilities and multi-line macros.

There are a few drawbacks to MICRO-C that the C purist might criticize. Currently there is no support for structures, unions, or bit-fields. There are also no floating point math functions or operations. For the most part, the effect of these omissions on an 8-bit embedded system is negligible. In most instances coding techniques can easily work around these limitations. Of course if these things really bother you there's always those $2,000 systems.

A word of caution--purchasing MICRO-C or any other high level language will not relieve you from learning the intricacies of the hardware for your target system. Unlike your PC or CP/M machine, there is no operating system or system services; more correctly, there are only those services you provide in the code you develop. There are no free rides when you're working with computers at this level.

Writing articles doesn't mean I'm the last word on the subject of microcontroller hardware and software development. There are lots of talented folks out there and I'm sure a lot of them are more experienced than I am. Do you have a favorite piece of software you use that you think could help others? Drop me a note and tell me about it. I'd especially like to hear from you if you're using something other than an MSDOS computer. Let me know what's available for the developer that uses an Amiga, Atari, Color Computer, or CP/M machine. When writing please remember to tell me what machine it runs on, who publishes it, and where to get it (BBS, direct sales, etc.) In fairness to all, please indicate if you are somehow associated with the publisher/manufacturer. I'll still mention your product if it's related to embedded systems work but I want to keep things on the up and up.

Author Information

Tim McDonough makes a living integrating remote measurement and control systems with data communications systems and developing hardware and software systems for process control and data acquisition. He prefers to be contacted electronically if possible.

Internet E-Mail: tmcdo450@athenanet.com
BIX: tmcdonough
US Mail: Tim McDonough
Suite 3-672
1405 Stevenson Drive
Springfield, IL 62703
USA

Vendors Mentioned in this article:

PseudoCorp (PseudoSam Assemblers/ Simulators) 716 Thimble Shoals Blvd. Suite E Newport News, VA 23606 (804) 873-1947

Dunfield Development Systems (MICRO-C/XASM) P.O. Box 31044 Nepean, Ontario (Canada) K2B 8S8 (613) 256-5820 (8:00am to 7:00pm EST)

CLASSIFIED, FOR SALE and SUPPORT WANTED

For Sale: GIMIX 6809 SS-50 floppy DMA disk controllers. Have six to sell at $25 each (plus shipping). These are like new and DOCS may be had for a fee. These were top of the line controllers for the SS-50bus system. Contact Bill at TCJ (916) 645-1670.

For Sale: GIMIX FULL SYSTEM, 6809 SS-50 box (very big and heavy), 15 slot mother board, with several SS-30 slots. 6809 CPU, 64K memory, 2 DMA disk controllers (8 and 5 1/4" drives supported), 2 serial ports, 1 720K 5 1/4" drive and 1 360K 5 1/4" drive in case. Books, Disks, works, some programs (Forth, wordprocessor, etc.),currently running FLEX. $100 (plus shipping if not in San Francisco bay area). Contact Neale at (415) 892-1432, or Bill at TCJ (916) 645-1670.

Support needed: Looking for help with South West Technical 6800 system. This is a tape cartridge unit (AC30). Trying to find tapes, docs, any help? Contact Todd Silk, at (208) 772,4631.

Wanted: SAGE/Stride CP/M68k or UNIX software. Have a SAGE IV, (actually II but labeled wrong) with 500K of memory. Also looking for hardrive for same. Want source code for CP/M BIOS and BOOT disk. have entire P-System with source code for that BIOS and ROMS. Complete set of DOCS and will copy what you need in exchange for software or support. Unix may not run on my system, but feel that I neeed to collect this information before it disapears for good. Interested in hearing about projects/uses you may be doing with these systems. Contact Bill at TCJ (916) 645-1670.

The Computer Journal Classified section is for items FOR SALE. it is priced and setup the same as *NUTS & VOLTS*. If you currently have an ad running in *NUTS & VOLTS*, just send us a copy of the ad and your invoice from them, along with your check, and we will publish it in the *TCJ*.

Classified ads are on a pre-paid basis only. The rate is $.30 per word for subscribers, and $.60 per word for others. There is a minimum $4.50 charge per insertion.

Support wanted is a free service to our readers who need to find old or missing documentation or software. No FOR SALE items allowed, however exchanges or like kind swapping is permitted. Please limit your request to one type of system. Call TCJ at (916) 645-1670.

Dr. S-100

By Herb R. Johnson

Regular Feature

Intermediate

Acquiring S-100 Skills

Resurrecting your S-100 system

My experience suggests that setting up a new S-100 system from old cards is like helping a person recover from a stroke: there is a period of tests to find damage, and stages of exercise to restore full functionality. New parts may be necessary, which also require an adjustment period. And, like a stroke, these events are rarely planned for.

I was rewriting a Tarbell floppy disk controller BIOS for a SD Systems floppy controller recently. My development system was a Vector Graphics S-100 chassis, an IMSAI front panel, a Cromemco Z80 ZPU, and four CCS 16K static RAM cards. The memory cards feature an LED which lights when the board is addressed, and DIP switch logic for independently addressing each 4K of memory. Static memory, by the way, is **strongly recommended** on the S-100 bus! Of course, the IMSAI front panel allows easy rebooting and memory examination. The Cromemco ZPU is switchable between 2 and 4 MHZ, ideal for debugging as you'll see later.

As is my custom, I ran the assembler and trotted to the "little programmers room." By the time I got back, all should have been assembled and the front panel blinking merrily away waiting for a keystroke. Instead, it was dark. After the appropriate calling down of the gods for daring to fail my "last" programming change of the day, my next reaction was..what? What do you do when your system fails?

Use your senses

Before you grab a scope or even a meter, **inhale through your nose repeatedly**, i.e. smell! Is something burning? With experience, you can easily tell the difference between the "toasty" smell of resistors, the "repaving the streets" smell of a fried transformer, and so on. Without those experiences, you can still trace the odor to some source. My luck, I noticed no such smells.

Next, **use your hearing**. Are the fans running? This will tell you if AC power is active (presuming you have AC fans. Some people, incidentally, have S-100 systems without fans. They call them "heaters.") Of course, **use your eyes!** Do you see smoke? Are some boards lit up (hopefully with LEDs, not with the incandescent glow of components). If you find any catastrophic damage, **shut down your system at once!**

Use your instruments

Total failure (no lights) is actually good news! That limits the problem to the power supply or the AC source. In my case, however, no such luck. The front panel lights were on, but not blinking. The next step was the trusty **voltmeter**. Mine is a Radio Shack digital voltmeter, cost about $25, but any meter will do. Measure the power supply voltages. Generally it is easiest to do from the enormous capacitors that old S-100 systems have: newer systems will have a nice set of power terminals to inspect. Measuring from the S-100 connector is **risky!!** Positive and negative voltages are adjacent: a brief short will at least ZAP your fuses, and possibly your logic!

Use your head

My system showed no +8 voltage: thus the S-100 cards had no source for their +5 volt regulators. However, the AC side of the power supply had 110V AC available. To test my system power supply, I removed most of the S-100 cards, both to minimize possible damage from any accident and to see if reduced power requirements (or disengaging a board's shorted regulator) would change the results.

In diagnosis there is a pattern of behavior to follow: **isolate, test, modify, repeat.** Isolate a subsystem, test its operation, modify a **single** feature of that subsystem, and test again. Repeat this cycle of modification and testing until you have either found a fault or verified the subsystem's operation. When you find a fault, **isolate** the fault area further and continue testing.

Forward and retreat

On my system, a resistance check of the rectifier diodes was good (forward resistance a few tens of ohms, reverse resistance very high). While a resistance check of the transformer did not suggest the windings were open, no voltage was produced at the secondary. What now?

In principle, I could have torn down the chassis and thoroughly tested all features of the power supply. In fact, I pulled out the rest of the S-100 cards, and went to the garage to get another S-100 chassis. Meanwhile, I commend Mr. Payson for his foresight in keeping several S-100 systems as a repair resource. If you, however, enjoy building power supplies or have a source of Vector Graphic transformers, you could have fixed my system. It's all a **matter of priorities.** Re-

member, mine was writing a new BIOS.

How to stabilize a "new" system

My new box was from a banking system, and had a full load of cards itself. There is an art to pulling out a dozen S-100 cards without shredding your hands on the IC pins which protrude like Punji sticks through each one. Hint: reach **below** each card to the bottom-most edge and **push up**. Practice helps.

From my previous use, I knew the power supply on this box was operational. It was certainly powerful enough, noting the number of chips on each card I just pulled out. To make space for my front panel card, however, I would need a **bus extender card**. (Why? The IMSAI front panel card is 19 inches long, and rests in a special slot on the bus that is extended outside the S-100 card cage, which is about 10 inches wide.) As my previous Vector Graphic box also required an extender, this was no problem. So, I plugged in the front panel, the processor card, one memory card and powered up the new system. The lights on all the boards came on! This suggested that no board was **seriously** damaged, at least.

"The helm is not responding, Captain..."

My first tests were, of course, to read memory locations via the front panel. I quickly found that, although I could read an address, I could not use the NEXT ADDRESS or DEPOSIT NEXT switches to proceed to the next memory address! Was this a failure of the processor card, or the front panel card? I know the answer, and decided to deal with it another day. (Remember the BIOS I was writing?) Who else out there knows the problem? I'll tell all next issue.

The fact that addressed locations had a variety of data values in them, as displayed from the front panel, suggested "nominal" operation of all cards. I then added the Tarbell controller, the serial I/O card and all memory cards, and attempted to boot. I guess I didn't mention that I was using two 8" floppy drives - the "small" half-height Mitibushi drives favored by Heath in their Z-100 systems. These drives, as do all 8" drives, have a distinctive **sound** to their operation that is **critical to diagnosing boot-up faults.** The sound of a "homing" drive is like drilling through wood, for instance.

Allow me to elaborate further, to demonstrate the value of "deep understanding" as well as the power of the senses to determine computer operations. The floppy boot process involves homing the drive (a seek to track 0), followed by a **click** or two as each track is sought (past tense of seek, sorry) and read. The contents of these "boot tracks" is the CP/M operating system including the BIOS specific to your machine. However, you may hear instead the "homing" sound followed by a sound like trying to drill through a metal plate and failing: that is due to the controller not sensing the **track 0** line from the drive to stop the "home" operation, probably because the **drive ready** line is apparently not active. In my case, this occurred because the floppy drive I used did not have "terminating resistors" on its control lines that give the interface driver chips a source of current sufficient to "pull down" the disk drive's **drive select** or **drive ready** lines. (See figure 1 for the schematic of this scenario.) Of course, I didn't know all this at once, so don't feel overwhelmed by this detail: experience comes with time and study.

More ailments

Once I added the terminating resistor pack, which was not needed by the Tarbell controller but necessary to the SD Systems controller, I heard the familiar "clunk-a-thunk" operation of the A drive. However, I did not get the final "A>" prompt on the screen. Again from experience, this is a typical symptom of some kind of **memory fault**. An area of memory vital to CP/M was absent, or a memory chip was bad, or the processor was running too fast for some area of memory to respond reliably.

Another possibility was **poor contact** on the S-100 bus by one or more cards. Unfortunately, this is a chronic problem in any system with hundreds of edge-contact points (i.e. number of cards multiplied by the number of bus connections per card). Age and storage make the problem more acute - kind of a "hardening of the arteries." The cure is a combination of edge connector cleaning and careful "seating" and reseating of the cards into the connectors. Observe any residue on the connectors. Try to use contact cleaner and a soft cloth to clean the edge of the board. Alcohol-based cleaners will degrade the plastics; and pencil erasers can be abrasive. A **soft** artist's eraser is better. Even pieces of paper, or a toothbrush with a **small** amount of soapy water followed by clear water, can make a difference. By the way, poor contacts are also a problem with **IC sockets:** examine these closely as well!

As for **memory tests**, try to run some large programs that will use all available memory; or switch suspect areas of RAM into or out of active memory space. A general memory test program is better still. You can at least use **CP/M's DDT** to fill and examine memory areas. Again, if your processor card has switchable speeds, run your system "too fast" to highlight errors. Note: **hot IC's**, including memory chips, **fail at lower speeds than cold chips.** For subtle errors, I've resorted to putting S-100 cards in the freezer and using a hair dryer to exaggerate temperature-based errors.

Running at last!

By a combination of the above, working on it over a few evenings, I finally re-established a running and stable system. All this may seem to be a lot of work to some of you. But, consider what we have learned about disk controllers, memory, contacts, and so on. The modular nature of the S-100 bus is ideal for repair and for self-education. My "lost" time is a gain for you if you have picked up some hints on testing and repairing **your** system! Have fun, and tell me what you have learned when it's your turn to resurrect your computer.

Letters and notes

Larry Cameron writes from GEnie (via his IMSAI!) on how he feels he may be "the only person crazy enough to be doing something like restoring old 'obsolete' computers" and hopes that TCJ

will "continue S-100 stuff." Larry, I hope my previous column shows a good beginning. I hear from a number of people who enjoy "antique computing," but rarely are any two of them in the same city! For communications, you have to go to national or international sources. I recommend the CP/M Tech echo on the BBS FidoNet network; I've also heard good reports about a similar echo on the Internet. Look at the end of my column for contacts.

Eliot Payson of Littleton CO says he "has several S-100 machines in various stages of disrepair." Well, join the club! Your condition is normal for the true S-100 fan! I think you'll like my column this month! Eliot also sent me, incidentally, a copy of his club's newsletter, and I hope our Editor Bill Kibler will "plug" it in more detail.

Some time ago, **Gary Stagliano** of Manchester CT wrote to me and asked about **MYZ80.** I'll refer him to the Sept/Oct 92 issue of TCJ for the brief from **Lee Bradley**. Incidentally, thanks to Lee's BBS, I have the latest version of **Simeon Cran's** MYZ80. As an early beta tester, I too am encouraged by this product. However, as currently distributed it lacks development documentationessential for any "tinkering" programming. I'll encourage such docs by contributing some of my "investigations" to Lee's board. See his column for contact info.

Resources

FidoNet: an international network of Bulletin Board Systems, usually with small or no subscription charges. It supports a number of informational exchange areas, including one for CP/M called "CP/M Tech." Contact your local computer club for a list of local boards which may carry FidoNet.

Internet: another international computer network, but harder to get onto for free as it is usually associated with universities or other institutions. Contact your local university's computer center for details.
Lee Bradley's Z-Node #12: (203) 665-1100 in CT.

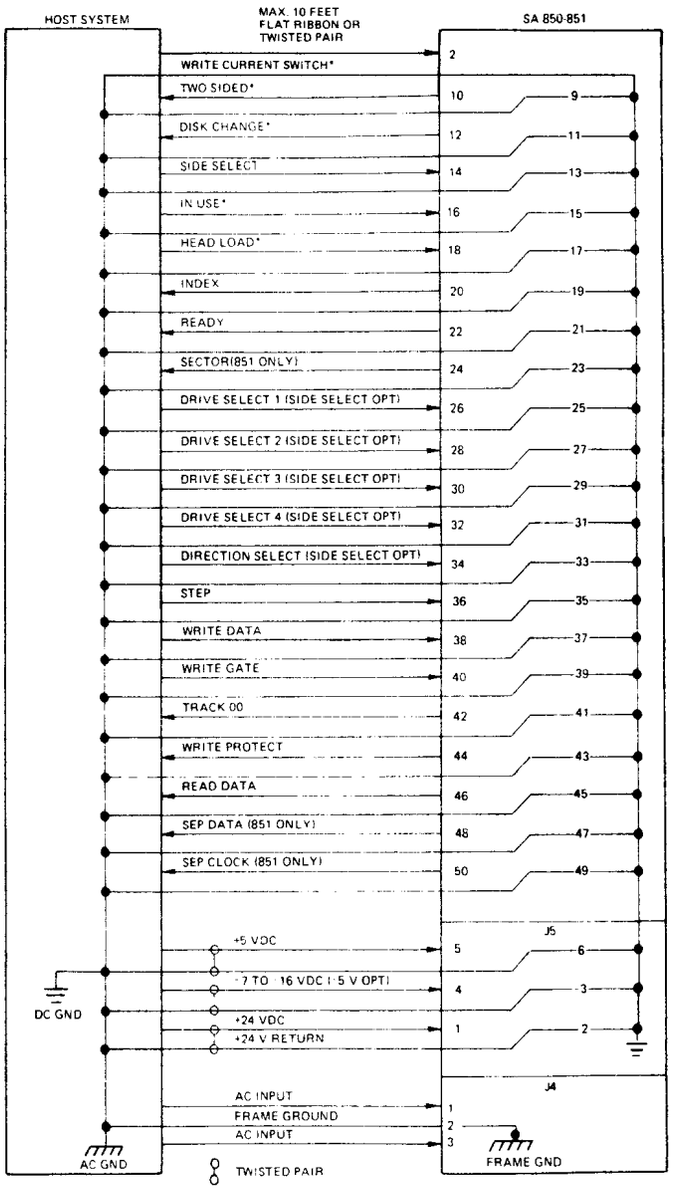

*These lines are alternate input/output lines and they are enabled by plugs. Reference section 7 for uses of these lines. Not shown are pins 4, 6, and 8 which are laternate I/O pins.

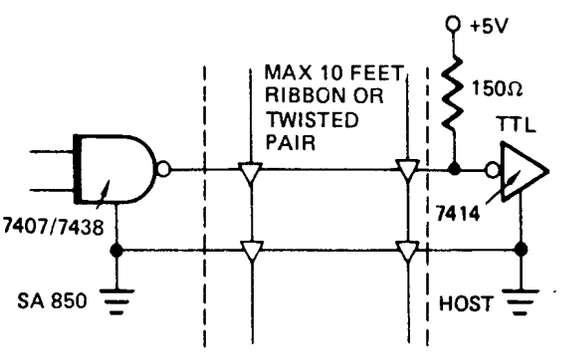

Figure 1: Disk Drive and Controller Interface

Regular Feature

Kaypro Support

Moving the Reset Button

Mr. Kaypro

By Charles B. Stafford

This is the first of a series of articles, the thrust of which is to make your hardware more friendly to the only user that counts, YOU. All of the modifications described will be possible with only ordinary hand tools, all of these modifications have been done by the author who has two left thumbs and is supposed (according to his mother) to be right-handed. Although the platform involved is a Kaypro, the principles are applicable to any machine, and the use of the Kaypro as an example, is only because that's what I happen to own. I am available for questions most evenings at (916) 482-8305.

The RESET Button

The Reset button on the Kaypro is mounted on the rear panel some- where in the vicinity of the top right corner. I say "somewhere" because the exact location varies with the age and type. Perhaps I am more adventurous than most, but I seem to get my trusty sidekick "hung" rather frequently, and it occurred to me once as I wished for a ten foot right arm, that I could correct Kaypro's design error. I'm sure that the designers put the Reset button where they did out of concern that it not be accidentally activated and to make sure that its use was the result of a conscious decision, but it certainly was inconvenient. The bottom line was that it belonged it on the front panel.

A little research (removing the cover) revealed that the proposed location was clear of any interference, and that the wires were long enough to reach. It also became apparent that the actual button was mounted from the rear of the panel, and held by a serrated ring screwed on from outside, and only finger tight.

Here is how we do it.

As it says in all the manuals, unplug the power cord not only from the wall, but also from the computer (unless you have a really early one where the line cord is mounted in a strain relief). Remove the cover carefully, and put the screws (10 for a Kaypro) in a safe place. I use those plastic cans that film comes in to keep track of the screws.

Decide where the new location is, make sure it will be clear of any folding keyboards, and leave some room for another front panel modification later. Mine is centered between the monitor and the drives, about an inch and a half above dead center vertically. Put a piece of masking tape over the spot, do your final measurements and mark the spot on the tape. If you have a center punch and are a purist, this is the time to use it.

Next comes the sneaky part. We use a six inch piece of two and a half inch duct tape, you know, McGyver's favorite tool, and make a tent over the back side of the new location (on the inside surface of the front panel). Pinch the sides of the tent shut and we have a "pouch" over the place we're going to drill to catch all the shavings and filings.

Here comes the scary part !

Using your favorite electric (OK, Hand) drill a 3/16" pilot hole being careful not to hit the duct tape, and then a 3/8" hole.

Remove the masking tape, smash the duct tape pouch down on the inside of the front panel to trap all the nasty things and remove it carefully along with the trash. Clean up the hole if necessary and the hard part is done.

One last bit of surgery, the wires going to the Reset button are fastened to the right stand-off (support) for the motherboard by a small plastic ty-wrap which must be cut. Diagonal cutters work just fine.

The fun part

Unscrew the plastic ring that secures the Reset button, move the button assembly to the new hole and reinstall the retaining ring.

Carefully check the area under the new location to make sure no metal chips got away, replace the cover and screws and

YOU'RE FINISHED.

When you realize the convenience of the new location, you may start thinking about the video brightness control. The only differences are 1. the knob ; 2. the wires will have to be extended. Otherwise the operations are the same.

Good Luck, and May You Live Long.

Chuck has proposed some articles for his next column unless our readers flood him with requests for information and help. The proposed articles include: A replacement of the 65 watt power supply with a pc-xt 150 watt power supply. A 5 Mhz speedup. A side select modification to convert early ssdd machines to dsdd. A Construction of a 4-drive decoder (TurboRom). If you have any special ideas or problems, give Chuck a call, and he will help you in resloving those Kaypro problems. BDK

Computing Timer Values

By Clem Pepper

Special Feature

Intermediate Support

Programs for Hardware

"It's test time", at least that's what I say to my students. For those who like to see if you remember your fundamental electronic skills, here is an article about timers and calculating the resistors and capacitors. This article also shows an appropriate use for doing something in "C." Read on and see if you can pass the fundamentals test. BDK

COMPUTE IC MONOSTABLE AND TIMER VALUES WITH THESE C PROGRAMS

IC monostable equations appear to be straightforward for the most part. They should present but few difficulties. But there is a catch: keeping track of the decimal points. Resistors are in M ohms, K ohms, or just plain ohms. Capacitance in micro-, nano-, and picofarads. Pulse widths in whatever submicrofraction of a second. Not to mention time spent thumbing through data manuals in search of an equation.

Well, here are two programs to do away with all that. The one, MONO.C, solves for any one of three variables when given the two known for thirteen popular IC monostables. The other, TIMR.C, solves for two of the five variables for the 555/556 timer in the astable mode. (Solutions for the monostable mode are provided in MONO.C.) Both programs are fully menu driven. In addition to calculation of component values and/or timing testing is performed for over/under limit conditions where these occur. If a limit is exceeded an explicit error message is displayed.

The programs are written in C. They should compile with any C compiler. I took care to avoid code that would limit the portability. For that reason MS DOS ANSI functions are used for cursor and screen clearing tasks. My original code was written in 1986 using The Software Toolworks TOOLWORKS C compiler. I recently recompiled it with Borland's Turbo C version 2.0 with only minor revisions. If you do not have a C compiler a disk containing both the source and executable code for a PC or compatible is available from the author.

Using The Monostable Program

For PC clone users, be sure that DEVICE=ANSI.SYS is included in your computer's CONFIG.SYS.

Table 1 is a listing of the timing equations for the IC monstables. In its operation the program first displays the 13 IC types. It then queries for the type and the variable to solve for. The three monostable variables are the pulse width, the timing resistance, and the timing capacitance. You are asked to enter the two known. The solution for the third is then calculated and displayed. If a component value exceeds a limit for the device an error message is displayed. The program can be run repeatedly without exiting from it.

The program begins with a display of the monostables for which a solution can be made. This is:

```
Hi! Enter the letter matching the mono of your choice
from the selection below.

A = 9600   E = 74122    I = 74LS221
B = 9601   F = 74123    J = 74C221/4538
C = 9602   G = 74LS123  K = 4528
D = 74121  H = 74221    L = 555/556
```

The first step, then, is to enter the letter corresponding to the device of interest.

In response the program prints a verification request using the device number.

```
Is the 74LS123 the mono you want? <Y/N> : y
```

Pressing "y" (or "Y") causes the program to continue. If "n" (or "N") is pressed you are asked to make another selection.

Assuming the entry is for the desired mono you are asked to select from one of the following options:

```
Type in T if the program is to solve for the pulse width:
Type in C if the program is to solve for the timing capacitor:
Type in R if the program is to solve for the timing resistor:
```

Suppose we enter "t" (or "T") for the pulse width. We are then queried for the two remaining requirements - the timing components:

```
Enter your timing capacitor requirement, CV : 1
Enter multiplier M=mfd : N=nfd : P=pfd : n
Enter your timing resistor requirement, RV : 47
Enter multiplier O=ohms : K=kohms : L=mohms : k
```

In this example we are using a one mfd capacitor with a 47K resistor. With this information provided the program calculates the pulse width and displays it's value along with the two given inputs.

TV = 21.15 uSec is the pulse width solution.
CV = 0.001 MFD is the timing capacitor solution.
RV = 47000 OHMs is the timing resistance solution.

We are then given the opportunity to continue with another or the same device with changed values or to exit the program.

Do you wish to continue with new input <Y/N>? : n

This sequence as it appears on our screen is illustrated in Figure 1.

The Monostable Program

The program source code is given in Listing 1. (The line numbers are for convenience only, not a part of the code.) Comments are used throughout to aid in its understanding. Defines and global declarations head the listing, preceding the instruction code. Program operation begins with the function **main()**. There is nothing particularly devilish about the code with the possible exception of the use of structure variables. Their primary purpose is to simplify the logic for the **printf()** statements.

Each function call is preceded by comment describing its purpose. A comment stating the source of the call is also provided. Each call to **scanf()** is followed by a call to flush the keyboard buffer. Without this call the program will leap over the next query. This arises from a bug in the **scanf()** library function.

The heart of the program is in the function **mon_calc()**. Here is where the program learns which variable it is to solve for. The variables ms, rn1, rn2, and rn3 (refer to Table 1) are then assigned. Variable md1 was assigned earlier in **mono_inp()** when you entered the device letter.

Following the assignments the program queries for the two known quantities and their multipliers. The actual multiplier is assigned with a call to **sca_par()**. A call is then made to **equatn()** for calculation of the unknown. The assignment variable **ms** is passed in the function call. Note that assignments for three quantities - TV, CV, RV - are made in this function.

On return from **equatn()** the three values are displayed. An advantage of C's structure is seen in the simplicity of the **printf()** statements.

A test then follows for limits. Should one or more be exceeded an error message is displayed.

Using The Timer Program

Table 2 is a listing of the 555/556 equations. In its operation the program first displays the five variables and instructions on their selection. You are asked to enter the two unknowns for which you want solutions. You are asked to verify these - a chance to change if desired. You must then provide information on the remaining three variables. Solutions are then calculated and displayed. No limits on component value exist. There is, however, an upper limit of 0.33 for the duty cycle - an error message is displayed if this is exceeded. Also for negative values in any solution. The program can be run repeatedly without exiting from it.

On typing TIMR and pressing RETURN the screen first clears and then displays the following:
Hi! Pick two variables for solution from the table below.
Enter them in ascending order (b e, not e b).
Follow each with a RETURN:

A = ONE CYCLE TIME, TT. TT = t(low)+t(high).
B = THE DUTY CYCLE, DC. DC = t(low)/TT
C = THE TIMING CAPACITANCE, TC.
D = THE UPPER RESISTANCE, RA.
E = THE LOWER RESISTANCE, RB.

Note: Do not select A and C as a pair.
An entry is not a commitment till you make it so.
If you flub, just continue on, till queried.
Then enter N to start over.

Note that we cannot obtain a solution for the duty cycle and the timing capacitance as a set.

In this example the solution for the two resistors, RA and RB, is desired, so we type "d" (or "D") followed by a RETURN, the "e" (or "E"). Again followed by a RETURN. Like so:

d
e

We are given an opportunity to change our minds:

Are D and E what you want? <Y/N> : y

Assuming happiness with our selection we enter "y" (or "Y"). We are then queried for the three remaining variables:

Enter your one-cycle time requirement TT: 5
Enter multiplier S=seconds : T=mS : U=uS : : t
Enter your duty cycle (must be <.33) requirement DC: .2
Enter multiplier - No multiplier for DC. Press RETURN.:
Enter your timing capacitor requirement TC: .05
Enter multiplier M=mfd : N=nfd : P=pfd : : m

In this we have chosen 5 mS for the period, a duty cycle of 0.2,

and a timing capacitance of .05 mfd. The resistance solutions follow:

RB = 2.89e+04 OHMS is the lower resistance solution.
RA = 8.66e+04 OHMS is the upper resistance solution.
TT = 5e+03 : DC = 0.2 : TC = 0.05

As with the monostable program we have the option of continuing:

Do you wish to continue with new input <Y/N>? n

Usually we know the frequency and duty cycle of the desired waveform. Capacitor values adhere closely to a fixed numerical scheme. The most flexibility and our greatest need to know is with the resistors. The effect of going to the nearest standard resistances is easily checked by their substitution and solving for the period, capacitance, and duty cycle.

The preceding sequence as it appears on our screen is illustrated in Figure 2.

The Timer Program

The program source code is given in Listing 2. It is similar in its construction to MONO.C. Although shorter it is somewhat more complex in having to deal with five variables. As with MONO.C the program takes advantage of #defines and the structure. The structure may look strange in that there are repetitions for the resistance variables RA and RB. This arises from the multiple solutions for these and the manner in which the data is displayed. Table 2 includes the program assignments associated with each of the possible solutions for all the variables.

The heart of the program is in the function **ast_calc()**. As in **mon_calc()** the program learns which variables relate to the unknowns (refer to Table 2). This function is more complex than its monostable counterpart in the number of possible combinations that exist. As in MONO.C scaling takes place in **sca_par()**.

Two calls are made to **equatn()** and the values displayed. A test for the duty cycle and any negative results follow.

Reference

Pepper, Clement S.
"A Monostable Catalog For Experimenters"
Popular Electronics, September, 1979, pp 69 - 79.

A Program Disk

A disk containing the equation tables, C source code, and the run files (.EXE) for the two programs is available from the author. Disk format may be any IBM compatible mode: 5 1/4 360 K/1.2M or 3 1/2. To obtain the disk send a check or money order for six dollars ($6.00) along with your mailing address to

C. S. Pepper,
13409 Midland Road, Apt. 175
Poway, CA 92064

A photocopy of the referenced article is also available for $2.00. Allow three to four weeks for delivery of these.

```
1: /* ********************************************** */
2: /* ****************** MONO.C ****************** */
3: /* ********************************************** */
4: /* A program for calculation of IC monostable  */
5: /* variables. ********************************* */
6: /* by Clement S. Pepper *********************** */
7: /* ********************************************** */
8:
9: #include <stdio.h>
10:
11:  /* == function prototypes == */
12: void main(),mono_inp(), error_in(), y_n_ver(),
13:  mon_calc(), equatn(int), sca_par(char), repeat();
14:
15: /*  DEVICE=ANSI.SYS required in CONFIG.SYS  */
16: #define CLRSN "\033[2J"  /* Clear the screen   */
17: #define CURUP "\033[A"   /* Cursor up one row  */
18: #define CURBK "\033[D"   /* Cursor left one col */
19: #define CURLS "\033[K"   /* Erase to line's end */
20: #define ERASE " "        /* Erase one space    */
21: #define POSCUR "\033[9;1H"
22:             /* position cursor line 9, col 1 */
23:  /* == global variables == */
24: int rerr;  /* Resistance limit error flag      */
25: int mdi;  /* Monostable device identifier      */
26: float sln; /* A monostable solution transfer var */
27: float sen; /* A monostable scaling transfer var  */
28: float se1; /* Transfer variable for 1st var     */
29: float se2; /* Transfer variable for 2nd var     */
30: float TV;  /* Timing variable (pulse width)     */
31: float CV;  /* Capacitance variable              */
32: float RV;  /* Resistance variable               */
33:
34: /* == mono identification structure == */
35: /* == Relates menu to device type == */
36: struct mono_sel {
37:       char *ic_sel;  /* User selection code  */
38:       char *ic_dev;  /* Monostable ID number  */
39:             } selattr[12] = {
40: /* ic_sel   ic_dev  */
41: /* ------  ------  */
42:  "A",   "9600",
43:  "B",   "9601",
44:  "C",   "9602",
45:  "D",   "74121",
46:  "E",   "74122",
47:  "F",   "74123",
48:  "G",   "74LS123",
49:  "H",   "74221",
50:  "I",   "74LS221",
51:  "J",   "74C221/4538",
52:  "K",   "4528",
53:  "L",   "555/556"
54:                       };
55:
56: /* == mono solution message structure == */
57: struct monocal {
58:  char *parmtr;  /* Parameter variable assigned  */
59:  char *p_strg;  /* Mono parameter string msg   */
60:  char *solmsg;  /* The solution message        */
61:  char *multi;   /* Data scaling factors        */
62:                } calattr[3] = {
63: /* parmtr  p_strg          solmsg          */
64: /* ------  ------          ------          */
65: /*  multi                                  */
66: /*  ----                                   */
67: "TV","pulse width","uSec is the pulse width solution.",
68:    "S=Sec : T=mS : U=uS",
69: "CV","timing capacitor",
    "MFD is the timing capacitor solution.",
70:    "M=mfd : N=nfd : P=pfd",
71: "RV","timing resistor",
    "OHMs is the timing resistance solution.",
72:    "O=ohms : K=kohms : L=mohms"
73:                          };
74:
75: /* == resistance error structure */
76: struct r_error {
77:   char *r_limit;
78:             } erattr[9] = {
```

```
79: "<1400","<2000","<5000","<10000",">40000",
80: ">50000",">100000",">350000",">1000000"
81:                 };
82: /* == Begin Program == */
83: void main()
84: {
85:   printf("%s",CLRSN);  /* Clear the screen */
86:
87:   /* **  Mono selection menu ** */
88:   printf("   Hi! Enter the letter matching the
            mono of your choice\n");
89:   printf("      from the selection below.\n\n");
90:   printf("      A = 9600   E = 74122   I = 74LS221 \n");
91:   printf("      B = 9601   F = 74123   J = 74C221/4538\n");
92:   printf("      C = 9602   G = 74LS123  K = 4528   \n");
93:   printf("      D = 74121  H = 74221   L = 555/556 \n\n");
94:    mono_inp();
95: }
96:
97: /* == Assign value to mono type identifier, mdi == */
98: void mono_inp()  /* From: main(), error_in(), y_n_ver() */
99: {
100:  char sel;
101:  sel = getch();
102:   switch(toupper(sel))  {
103:    case 'A': mdi = 0;  y_n_ver(); mon_calc(); /* 9600      */
104:    case 'B': mdi = 1;  y_n_ver(); mon_calc(); /* 9601      */
105:    case 'C': mdi = 2;  y_n_ver(); mon_calc(); /* 9602      */
106:    case 'D': mdi = 3;  y_n_ver(); mon_calc(); /* 74121     */
107:    case 'E': mdi = 4;  y_n_ver(); mon_calc(); /* 74122     */
108:    case 'F': mdi = 5;  y_n_ver(); mon_calc(); /* 74123     */
109:    case 'G': mdi = 6;  y_n_ver(); mon_calc(); /* 74LS123   */
110:    case 'H': mdi = 7;  y_n_ver(); mon_calc(); /* 74221     */
111:    case 'I': mdi = 8;  y_n_ver(); mon_calc(); /* 74LS221   */
112:    case 'J': mdi = 9;  y_n_ver(); mon_calc(); /* 74C221/4538 */
113:    case 'K': mdi = 10; y_n_ver(); mon_calc(); /* 4528      */
114:    case 'L': mdi = 11; y_n_ver(); mon_calc(); /* 555/556   */
115:     default: printf("%s%s%s%s%s",CURUP,CURBK,ERASE,CURBK);
116:      error_in();
117:                }
118: }
119:
120: /* == Display error message in menu selection == */
121: void error_in()  /* From: mono_inp() */
122: {
123:   printf("      Your entry is outside the menu.\n");
124:   printf("      Please re-enter your selection.\n");
125:    mono_inp();
126: }
127:
128: /* == Requestor selection confirmation == */
129: void y_n_ver()  /* From: mono_inp() */
130: {
131: char verify;
132:    printf("\tIs the %s the mono you want? <Y/N> : ",
           selattr[mdi].ic_dev);
133:    verify = getch();
134:     if(toupper(verify) == 'Y') { printf("\n"); return; }
135:    else printf("%s%s",POSCUR,CURLS);
136:        printf("\tPlease re-enter your selection.\n");
137:         mono_inp();
138: }
139:
140: /* == Variable inputs and calculations; display of results. == */
141: void mon_calc()  /* From: mono_inp() */
142: {
143: char par, scale;
144:  /* par - parameter input for solution: T, C, or R */
145:  /* scale - multiplier input by user            */
146: int m1, m2, m3, ms;
147:  /* m1, m2, m3 - struct monocal routing code   */
148:  /* ms - monocal structure identifier          */
149:
150: printf("\tType in T if the program is to solve for the
         pulse width: \n");
151: printf("\tType in C if the program is to solve for the
         timing capacitor: \n");
153:  printf("\tType in R if the program is to solve for the
         timing resistor: \n");
155: /* **  Enter parameter for solution: 'R', 'C', or 'T' ** */
156:  par = getch();
157:
158: /* ** Assign variables for identifying parameters R, C, or T ** */
159:
160:  /* ** solve for the resistance ** */
161:      if(toupper(par) == 'R') {
162:        m1 = 0; m2 = 1; m3 = 2;
163:      if(mdi == 0 || mdi == 1 || mdi == 4)  { ms = 17; }
164:        /* 9600, 9601, 74122 */
165:   else if(mdi == 2)  {  ms = 18; } /* 9602 */
166:   else if(mdi == 3 || mdi == 7 || mdi == 8)  { ms = 19; }
167:        /* 74121, 74221, 74LS221 */
168:   else if(mdi == 5)  {  ms = 20; } /* 74123 */
169:   else if(mdi == 6)  {  ms = 21; } /* 74LS123 */
170:   else if(mdi == 9)  {  ms = 22; } /* 74C221,4538 */
171:   else if(mdi == 10) {  ms = 23; } /* 4528 */
172:   else if(mdi == 11) {  ms = 24; } /* 555/556 */
173:                      }
174:  /* ** solve for the pulse width ** */
175:   else if(toupper(par) == 'T') {
176:         m1 = 1; m2 = 2; m3 = 0;
177:      if(mdi == 0 || mdi == 1 || mdi == 4)  {  ms =  1; }
178:        /* 9600, 9601, 74122 */
179:   else if(mdi == 2)  {  ms =  2; } /* 9602 */
180:   else if(mdi == 3 || mdi == 7 || mdi == 8)  {  ms =  3; }
181:        /* 74121, 74221, 74LS221 */
182:   else if(mdi == 5)  {  ms =  4; } /* 74123 */
183:   else if(mdi == 6)  {  ms =  5; } /* 74LS123 */
184:   else if(mdi == 9)  {  ms =  6; } /* 74C221,4538 */
185:   else if(mdi == 10) {  ms =  7; } /* 4528 */
186:   else if(mdi == 11) {  ms =  8; } /* 555/556 */
187:                      }
188:  /* ** solve for the capacitance ** */
189:   else if(toupper(par) == 'C') {
190:         m1 = 2; m2 = 0; m3 = 1;
191:      if(mdi == 0 || mdi == 1 || mdi == 4)  {  ms =  9; }
192:        /* 9600, 9601, 74122 */
193:   else if(mdi == 2)  {  ms = 10; } /* 9602 */
194:   else if(mdi == 3 || mdi == 7 || mdi == 8)  {  ms = 11; }
195:        /* 74121, 74221, 74LS221 */
196:   else if(mdi == 5)  {  ms = 12; } /* 74123 */
197:   else if(mdi == 6)  {  ms = 13; } /* 74LS123 */
198:   else if(mdi == 9)  {  ms = 14; } /* 74C221,4538 */
199:   else if(mdi == 10) {  ms = 15; } /* 4528 */
200:   else if(mdi == 11) {  ms = 16; } /* 555/556 */
201:                      }
202:
203:  /* ** Enter 1st known parm and scaling required ** */
204:   printf("\n");
205:   printf("\tEnter your %s requirement, %s : ",calattr[m1].p_strg,\
206:          calattr[m1].parmtr);
207:    scanf("%f",&se1); fflush(stdin);
208:
209:   printf("\tEnter multiplier %s : ",calattr[m1].multi);
210:    scanf("%c",&scale);  fflush(stdin);
211:     sen = se1; sca_par(scale); se1 = sen;
212:
213:  /* ** Enter 2nd known par and required scaling ** */
214:   printf("\tEnter your %s requirement, %s : ",calattr[m2].p_strg,\
215:          calattr[m2].parmtr);
216:    scanf("%f",&se2);  fflush(stdin);
217:
218:   printf("\tEnter multiplier %s : ",calattr[m2].multi);
219:    scanf("%c",&scale);  fflush(stdin);
220:     sen = se2; sca_par(scale); se2 = sen;
221:
222:  /* ** Get solution for unknown parameter ** */
223: equatn(ms);
224:  printf("\n\t%s = %g %s\n",calattr[m3].parmtr,sln,calattr[m3].solmsg);
225:  printf("\t%s = %g %s\n",calattr[m1].parmtr,se1,calattr[m1].solmsg),
226:  printf("\t%s = %g %s\n",calattr[m2].parmtr,se2,calattr[m2].solmsg);
227:
228:       /********** ERROR MESSAGES **********/
229:     /* ** RESISTANCE UNDER/OVER LIMIT VALUES ** */
230:   if(mdi == 3)  { if(RV <1400) rerr = 1;  }
231:   if(mdi == 7 || mdi == 8) { if(RV <2000) rerr = 2;  }
232:   if(mdi == 0 || mdi == 1 || mdi == 2 || mdi == 4  || mdi == 5
233:            || mdi == 6 || mdi == 10 )
234:          { if(RV <5000) rerr = 3;  }
235:   if(mdi == 0 || mdi == 1 || mdi == 2 ) { if(RV >25000)  rerr = 4;  }
236:
237:   if(mdi == 3)  { if(RV >40000)  rerr = 5;  }
238:   if(mdi == 4 || mdi == 5 || mdi == 6)  { if(RV >50000)  rerr = 6;  }
239:   if(mdi == 8)  { if(RV >100000)  rerr = 7;  }
240:   if(mdi == 9)  { if(RV >350000)  rerr = 8;  }
241:   if(mdi == 10) { if(RV >1000000) rerr = 9;  }
242:
243:  /* ** Print resistance under/over error message ** */
244:     if(rerr != 0)       {
245:      printf("\n\tA resistance %s is unacceptable for the %s.\n",\
246:           erattr[rerr-1].r_limit,selattr[mdi].ic_dev); rerr = 0;
247:       printf("\tSelect a new value for R and begin again.\n");
248:                 }
249:     /* ** CAPACITANCE OVER LIMIT VALUES ** */
250:     if(mdi == 3 || mdi == 7 || mdi == 8)  { if(CV >1000)  {
251:    printf("\n\tCapacitance >1000 mfd is unsatisfactory for the %s .",\
252:        selattr[mdi].ic_dev);
253:    printf("\n\tSelect new a new value for C and begin again."); }
254:                                          }
255:   repeat();
256: }
257:
258: /* == Solve for the unknown parameter == */
259: void equatn(ms)  /* From: mon_calc() */
260: int ms;
261: {
262: int lvcc;  char dcv;
263:  /* lvcc - natural log value for VCC    */
264:  /* dcv - code for 4528 DC voltage input */
265:     if(ms == 7 || ms == 15 || ms == 23)  {
266:  printf("\n\tEnter one of the following code numbers for the 4528 VCC");
267:  printf("\n\t  If your VCC is 5 volts enter 1");
```

```
printf("\n\t  If your VCC is 10 volts enter 2");
printf("\n\t  If your VCC is 12 volts enter 3");
printf("\n\t  If your VCC is 15 volts enter 4");

/* ** input code for VCC ** */
  dcv = getch();

/* ** Calculate natural log of VCC ** */
    if(dcv == '1') lvcc = 1.609;
else if(dcv == '2') lvcc = 2.303;
else if(dcv == '3') lvcc = 2.485;
else if(dcv == '4') lvcc = 2.708;       }

/* ** Solve for the Pulse Width (TV) solution ** */
   if(ms == 1) { CV=se1; RV=se2; sln=.32*RV*CV*(1+.7/RV); TV=sln; }
else if(ms == 2) { CV=se1; RV=se2; sln=.31*RV*CV*(1+1/RV); TV=sln; }
else if(ms == 3) { CV=se1; RV=se2; sln=.693*RV*CV;       TV=sln; }
else if(ms == 4) { CV=se1; RV=se2; sln=.28*RV*CV*(1+.7/RV); TV=sln; }
else if(ms == 5) { CV=se1; RV=se2; sln=.45*RV*CV;        TV=sln; }
else if(ms == 6) { CV=se1; RV=se2; sln= RV*CV;           TV=sln; }
else if(ms == 7) { CV=se1; RV=se2; sln=.2*RV*CV*lvcc;    TV=sln; }
else if(ms == 8) { CV=se1; RV=se2; sln=1.1*RV*CV;        TV=sln; }

/* ** Solve for the Capacitance (CV) solution ** */
else if(ms == 9) { RV=se1; TV=se2; sln=3.125*TV/RV*(1+1/RV); CV=sln; }
else if(ms == 10) { RV=se1; TV=se2; sln=3.22*TV/RV*(1+1/RV); CV=sln; }
else if(ms == 11) { RV=se1; TV=se2; sln=1.443*TV/RV;       CV=sln; }
else if(ms == 12) { RV=se1; TV=se2; sln=3.57*TV/RV*(1+.7/RV); CV=sln; }
else if(ms == 13) { RV=se1; TV=se2; sln=2.22*TV/RV;        CV=sln; }
else if(ms == 14) { RV=se1; TV=se2; sln=TV/RV;           CV=sln; }
else if(ms == 15) { RV=se1; TV=se2; sln=5*TV/RV*lvcc;      CV=sln; }
else if(ms == 16) { RV=se1; TV=se2; sln=.909*TV/RV;        CV=sln; }

/* ** Solve for the Resistance (RV) solution ** */
else if(ms == 17) { TV=se1; CV=se2; sln=3.125*TV/CV;       RV=sln; }
else if(ms == 18) { TV=se1; CV=se2; sln=3.226*TV/CV;       RV=sln; }
else if(ms == 19) { TV=se1; CV=se2; sln=1.443*TV/CV;       RV=sln; }
else if(ms == 20) { TV=se1; CV=se2; sln=3.57*TV/CV;        RV=sln; }
else if(ms == 21) { TV=se1; CV=se2; sln=2.22*TV/CV;        RV=sln; }
else if(ms == 22) { TV=se1; CV=se2; sln=TV/CV;           RV=sln; }
else if(ms == 23) { TV=se1; CV=se2; sln=5*TV/CV*lvcc;      RV=sln; }
else if(ms == 24) { TV=se1; CV=se2; sln=.909*TV/CV;        RV=sln; }
}

/* == Multiply input by scale factor == */
void sca_par(multi) /* From: mon_calc() */
char multi;
{
    if(toupper(multi) == 'S') { sen=sen*1e+6; return; }
  else if(toupper(multi) == 'T') { sen=sen*1e+3; return; }
  else if(toupper(multi) == 'U') { sen=sen*1;   return; }
  else if(toupper(multi) == 'M') { sen=sen*1;   return; }
  else if(toupper(multi) == 'N') { sen=sen*1e-3; return; }
  else if(toupper(multi) == 'P') { sen=sen*1e-6; return; }
  else if(toupper(multi) == 'O') { sen=sen*1;   return; }
  else if(toupper(multi) == 'K') { sen=sen*1e+3; return; }
  else if(toupper(multi) == 'L') { sen=sen*1e+6; return; }
}

/* == Request decision to continue or to exit. == */
void repeat() /* From: mon_calc() */
{
char verify;
 printf("\n\tDo you wish to continue with new input <Y/N>? : ");
 verify = getch();
 if(toupper(verify) == 'N') exit();
 else printf("%s",CLRSN); main();
}
```

Listing 1. The source code for MONO.C.

```
/* ********************************************** */
/* ****************** TIMR.C ******************** */
/* ********************************************** */
/* A program for calculation of 555/556 timer ** */
/* variables. ******************************** */
/* by Clement S. Pepper ********************** */
/* ********************************************** */

#include <stdio.h>

/* == DEVICE=ANSI.SYS required in CONFIG.SYS == */
#define CLRSN "\033[2J" /* ANSI Clear the screen    */
#define CURUP "\033[A"  /* ANSI Cursor up  one line */
#define CURDN "\033[B"  /* ANSI Cursor down one col */
#define CURBK "\033[D"  /* ANSI Cursor left one col */
#define CURLS "\033[K"  /* ANSI Erase to line's end */
#define ERASE " "       /* Erase one space          */

/* == function prototypes == */
void main(), error_in(), astbl_pr(), astblin(),
   y_n_ver(), ast_calc(int,int), repeat(),
   sca_par(int), equatn(int), new_ent(int,int,int);

/* == global variables == */
int aj;                 /* xfer variable 2nd in */
int flg = 0;            /* 2nd entry flag       */
float sln,TT,DC,TC,RA,RB;  /* solution variables   */
float sen, se1, se2, se3;  /* known var's entered  */

/* == astable solution message structure == */
struct astcalc {
  char *eq_cde;  /* This letter and          */
  char *p_strng; /* this message together in */
  char *parm;    /* concert relate this variable */
  char *msg;     /* to this solution for unknown */
  char *multi;   /* Input parm multiplier.   */
           } astable[8] = {
  "A", "one-cycle time", "TT",
  "uSec is the solution for one period.",
  " S=seconds : T=mS : U=uS : ",
  "B", "duty cycle (must be <.33)", "DC",
  "is the solution for the duty cycle.",
  " - No multiplier for DC. Press RETURN.",
  "C", "timing capacitor","TC",
  "MFD is the timing capacitor solution",
  " M=mfd : N=nfd : P=pfd : ",
  "D", "upper resistor", "RA",
  "OHMS is the upper resistance solution.",
  " O=Ohms : K=Kohms : L=Mohms : ",
  "E", "lower resistor", "RB",
  "OHMS is the lower resistance solution.",
  " O=Ohms : K=Kohms : L=Mohms : ",
  "D", "upper resistor", "RA",
  "OHMS is the upper resistance solution.",
  " O=Ohms : K=Kohms : L=Mohms : ",
  "E", "lower resistor", "RB",
  "OHMS is the lower resistance solution.",
  " O=Ohms : K=Kohms : L=Mohms : ",
  "E", "lower resistor", "RB",
  "OHMS is the lower resistance solution.",
  " O=Ohms : K=Kohms : L=Mohms : " };

/* == Begin Program == */
void main()
{
  printf("%s",CLRSN); /* clear the screen */
  printf("\tHi! Pick two variables for solution from the
        table below.\n");
  printf("\tEnter them in ascending order (b e, not e b).\n");
  printf("\tFollow each with a RETURN:\n\n");
  printf("\t A = ONE CYCLE TIME, TT. TT = t(low)+t(high).\n");
  printf("\t B = THE DUTY CYCLE, DC. DC = t(low)/TT\n");
  printf("\t C = THE TIMING CAPACITANCE, TC.\n");
  printf("\t D = THE UPPER RESISTANCE, RA.\n");
  printf("\t E = THE LOWER RESISTANCE, RB.\n\n");
  printf("\tNote: Do not select A and C as a pair.\n");
  printf("\tAn entry is not a committment till you make it so.\n");
  printf("\tIf you flub, just continue on, till queried.\n");
  printf("\tThen enter N to start over.\n");
   astblin();
}

/* == erroneous input selection message == */
void error_in() /* from astblin() */
{
 printf("\n\tYour entry, %c, is outside the menu.\n",astable[aj].eq_cde);
  printf("\t Please re-enter your selection.\n");
  astblin();
}

/* == defining astable parameters == */
void astbl_pr() /* from astblin() */
{
int ast1, ast2;
     if(flg == 0) ast1 = aj;
      else ast2 = aj;

      if(flg == 0) { flg = 1; astblin(); }
      if(flg == 1) flg = 0;

printf("\tAre %s and %s what you want? <Y/N> :",astable[ast1].eq_cde,\
        astable[ast2].eq_cde);
    printf("\n");
    y_n_ver();
     ast_calc(ast1,ast2);
}

/* == read keyboard inputs == */
/* from main(), error_in(), astbl_pr(), y_n_ver(), ast_calc() */
void astblin()
{
char sel; scanf("%c",&sel); fflush(stdin);
   switch(toupper(sel)) {
      case 'A': aj = 0; astbl_pr();
      case 'B': aj = 1; astbl_pr();
      case 'C': aj = 2; astbl_pr();
      case 'D': aj = 3; astbl_pr();
      case 'E': aj = 4; astbl_pr();
       default: printf("%s%s%s%s",CURUP,CURBK,ERASE,CURBK);
         error_in();
```

```
120:                    }
121: }
122:
123: /* == verify selection entry == */
124: void y_n_ver() /* from astbl_pr() */
125: {
126: char verify;
127:    verify = getch();
128:   if(toupper(verify) == 'Y') return;
129:  else {
130:    printf("%s%s%s%s%s",CURBK,CURUP,CURUP,CURUP,CURUP);
131:    printf("%s\n%s\n%s\n%s\n",CURLS,CURLS,CURLS,CURLS);
132:    printf("%s%s%s%s%s",CURBK,CURUP,CURUP,CURUP,CURUP);
133:    printf("\t  Please re-enter your selections.\n");
134:
135:     flg = 0;  astblin();
136:     }
137: }
138:
139: /* == perform calculations == */
140: void ast_calc(n1,n2) /* from astbl_pr() */
141:  int n1, n2;
142:  {
143: char scale, sf1, sf2;
144: int m1, m2, m3, ms1, ms2;
145:
146:    if(n1 == 0 && n2 == 1) { m1=2; m2=3; m3=4; ms1=0; ms2=1; }
147: else if(n1 == 0 && n2 == 2) {
148:  printf("\tThere is no unique solution for this combination.\n");
149:  printf("\tPlease re-enter your selections.\n");  astblin();  }
150: else if(n1 == 0 && n2 == 3) { m1=1; m2=2; m3=4; ms1=3; ms2=0; }
151: else if(n1 == 0 && n2 == 4) { m1=1; m2=2; m3=3; ms1=4; ms2=0; }
152: else if(n1 == 1 && n2 == 2) { m1=0; m2=3; m3=4; ms1=2; ms2=1; }
153: else if(n1 == 1 && n2 == 3) { m1=0; m2=2; m3=4; ms1=5; ms2=1; }
154: else if(n1 == 1 && n2 == 4) { m1=0; m2=2; m3=3; ms1=7; ms2=1; }
155: else if(n1 == 2 && n2 == 3) { m1=0; m2=1; m3=4; ms1=3; ms2=2; }
156: else if(n1 == 2 && n2 == 4) { m1=0; m2=1; m3=3; ms1=4; ms2=2; }
157: else if(n1 == 3 && n2 == 4) { m1=0; m2=1; m3=2; ms1=6; ms2=3; }
158:
159:  /* ** First parameter entry. ** */
160: printf("Enter your %s requirement %s: ",astable[m1].p_strng,\
161: astable[m1].parm);
162:  scanf("%f",&se1); fflush(stdin),
163: printf("Enter multiplier %s ",astable[m1].multi);
164:  scanf("%c",&scale);  fflush(stdin);
165:  sen = se1;  sca_par(scale); se1 = sen;
166:
167:  /* ** Second parameter entry. ** */
168: printf("Enter your %s requirement %s: ",astable[m2].p_strng,\
169: astable[m2].parm);
170:  scanf("%f",&se2);  fflush(stdin);
171:
172: printf("Enter multiplier %s ",astable[m2].multi);
173:  scanf("%c",&scale);  fflush(stdin);
174:  sen = se2;  sca_par(scale); se2 = sen;
175:
176:  /* ** Third parameter entry. ** */
177: printf("Enter your %s requirement %s: ",astable[m3].p_strng,\
178: astable[m3].parm);
179:  scanf("%f",&se3);  fflush(stdin);
180:
181: printf("Enter multiplier %s ",astable[m3].multi);
182:  scanf("%c",&scale);  fflush(stdin);
183:  sen = se3;  sca_par(scale); se3 = sen;
184:
185:    if(n1 == 0 && n2 == 1) { TC=se1;RA=se2;RB=se3; }
186: else if(n1 == 0 && n2 == 3) { DC=se1;TC=se2;RB=se3; }
187: else if(n1 == 0 && n2 == 4) { DC=se1;TC=se2;RA=se3; }
188: else if(n1 == 1 && n2 == 2) { TT=se1;RA=se2;RB=se3; }
189: else if(n1 == 1 && n2 == 3) { TT=se1;TC=se2;RB=se3; }
190: else if(n1 == 1 && n2 == 4) { TT=se1;TC=se2;RA=se3; }
191: else if(n1 == 2 && n2 == 3) { TT=se1;DC=se2;RB=se3; }
192: else if(n1 == 2 && n2 == 4) { TT=se1;DC=se2,RA=se3; }
193: else if(n1 == 3 && n2 == 4) { TT=se1,DC=se2;TC=se3; }
194:
195:
196:  /* ** display solutions ** */
197:  equatn(ms1);
198:  printf("\t%s = %.3g %s\n",astable[ms1].parm,sln,astable[ms1].msg);
199:  equatn(ms2);
200:  printf("\t%s = %.3g %s\n",astable[ms2].parm,sln,astable[ms2].msg);
201:
202:  printf("\t%s = %.3g : %s = %.3g : %s = %.3g \n",astable[m1].parm,se1,\
203:  astable[m2].parm,se2,astable[m3].parm,se3);
204:
205:        /* ********** ERROR MESSAGES ********** */
206:
207: if(DC > .333) {
208:  printf("\t**** DC must be less than 0.333. ****\n");
209:   new_ent(m1,m2,m3);  }
210:
211: if(TT<0 || DC<0 || TC<0 || RA<0 || RB<0)  {
212: printf("\t*** Negative values for any parameter not acceptable. ***\n");
213:   new_ent(m1,m2,m3);  }
214:
215:  printf("\n\tDo you wish to continue with new input <Y/N>?");
216:   repeat();
217: }
218:
219: /* == repeat option == */
220: void repeat() /* from ast_calc() */
221: {
222:  char verify;
223:   verify = getch();
224:   if(toupper(verify) == 'N')  exit();
225:    main();
226: }
227:
228: /* == parameter scaling == */
229: void sca_par(multi) /* from ast_calc() */
230: char multi;
231: {
232:     if(toupper(multi) == 'S') { sen = sen*1e+6;  return;  }
233:  else if(toupper(multi) == 'T') { sen = sen*1e+3;  return;  }
234:  else if(toupper(multi) == 'U') { sen = sen*1;     return;  }
235:
236:  else if(toupper(multi) == 'M') { sen = sen*1;     return;  }
237:  else if(toupper(multi) == 'N') { sen = sen*1e-3;  return;  }
238:  else if(toupper(multi) == 'P') { sen = sen*1e-6;  return;  }
239:
240:  else if(toupper(multi) == 'O') { sen = sen*1;     return;  }
241:  else if(toupper(multi) == 'K') { sen = sen*1e+3,  return;  }
242:  else if(toupper(multi) == 'L') { sen = sen*1e+6;  return;  }
243: }
244:
245: /* == solution calculations == */
246: void equatn(eqn) /* from ast_calc() */
247: int eqn;
248: {
249:  if(eqn== 0 ) { sln = .693*(RA+2*RB)*TC;       TT = sln; return; }
250:  if(eqn== 1 ) { sln = RB/(RA+2*RB);            DC = sln; return, }
251:  if(eqn== 2 ) { sln = 1.443*TT/(RA+2*RB);      TC = sln; return; }
252:  if(eqn== 3 ) { sln = RB*(1-2*DC)/DC;          RA = sln; return; }
253:  if(eqn== 4 ) { sln = DC*RA/(1-2*DC);          RB = sln; return; }
254:  if(eqn== 5 ) { sln = (TT-1.368*RB*TC)/(.693*TC); RA = sln; return; }
255:  if(eqn== 6 ) { sln = TT*DC/(.693*TC);         RB = sln; return; }
256:  if(eqn== 7 ) { sln = (TT-.693*RA*TC)/(1.386*TC); RB = sln; return; }
257: }
258:
259: /* == new parameter selections == */
260: void new_ent(m1,m2,m3) /* from ast_calc() */
261:  int m1,m2,m3;
262:  {
263:   printf("\tSelect new values for %s, %s, and %s for calculation.\n",\
264:    astable[m1].parm,astable[m2].parm,astable[m3].parm);
265: }
```

Listing. 2. The source code for TIMR.C.

IC MONO EQUATIONS AND VARIABLES USED IN MONO.C

PROGRAM VARIABLES

MENU	MONO	EQUATIONS	mdi	ms	m1	m2	m3	SOURCE/LIMITS
A	9600	R = 3.125*T/C*(1+.7/R)	0	17	0	1	2	Fairchild
		T = .32*R*C*(1+0.7/R)	0	1	1	2	0	R: 5 - 25K
		C = 3.125*T/(R*(1+.7/R))	0	9	2	0	1	C: ANY
B	9601	R = 3.125*T/C*(1+.7/R)	1	17	0	1	2	National
		T = .32*R*C*(1+0.7/R)	1	1	1	2	0	R: 5 - 25K
		C = 3.125*T/(R*(1+.7/R)	1	9	2	0	1	C: ANY
C	9602	R = 3.226*T/C*(1+1/R)	2	18	0	1	2	National
		T = .31*R*C*(1+1/R)	2	2	1	2	0	R: 5 - 25K
		C = 3.22T/R*(1+1/R)	2	10	2	0	1	C: ANY
D	74121	R = 1.443*T/C	3	19	0	1	2	Texas Inst.
		T = .693*R*C	3	3	1	2	0	R: 1.4 - 40K
		C = 1.443*T/R	3	11	2	0	1	C: ANY
E	74122	R = 3.125*T/C*(1+.7/R)	4	17	0	1	2	Texas Inst.
		T = .32*R*C*(1+.7/R)	4	1	1	2	0	R: 5 - 50K
		C = 3.125*T/R*(1+.7/R)	4	9	2	0	1	C: ANY
F	74123	R = 3.57*T/C*(1+.7/R)	5	20	0	1	2	Texas Inst.
		T = .28*R*C*(1+.7/R)	5	4	1	2	0	R: 5 - 50K
		C = 3.57T/R*(1+.7/R)	5	12	2	0	1	C: ANY
G	74LS123	R = 2.22*T/C	6	21	0	1	2	Texas Inst.
		T = .45*R*C	6	5	1	2	0	R: 5 - 50K
		C = 2.22*T/R	6	13	2	0	1	C: ANY
H	74221	R = 1.443*T/C	7	19	0	1	2	Texas Inst.
		T = .693*R*C	7	3	1	2	0	R: 2 - 40K
		C = 1.443*T/R	7	11	2	0	1	C: 1000 MFD
I	74LS221	R = 1.443*T/C	8	19	0	1	2	Texas Inst.
		T = .693*R*C	8	3	1	2	0	R: 2 - 100K
		C = 1.443*T/R	8	11	2	0	1	C: 1000 MFD

J 74C221	R = T/C	9 22 0 1 2	National
	T = R*C	9 6 1 2 0	R: ANY
	C = T/R	14 2 0 1	C: ANY
also the	4538 Motorola, RCA		
K 4528	R = 5*T/(C*(ln Vcc))	10 23 0 1 2	Motorola,RCA
	T = .2*R*C*(ln Vcc)	10 7 1 2 0	R: 5K - 1M
	C = 5*T/R*(ln Vcc)	10 15 2 0 1	C: ANY
	ln 5 = 1.61		
	ln 10 = 2.30		
	ln 12 = 2.48		
	ln 15 = 2.71		
L 555/556 Pulse High Time:			
	R = .909*T/C	11 24 0 1 2	Signetics
	T = 1.1*R*C	11 8 1 2 0	R: ANY
	C = .909*T/R	11 16 2 0 1	C: ANY

Note: Sources shown are for reference; each monostable type is typically available from several sources.

Table 1. Monostable equations as employed in MONO.C.

555/556 ASTABLE TIMER EQUATIONS

Period:	TT = .693*(RA+2*RB)*TC
Duty Cycle:	DC = RB/(RA+2*RB)
Timing Capacitor:	TC = 1.443*TT/(RA+2*RB)
Upper Resistor:	RA1 = (TT-1.368*RB*TC)/(.693*TC)
or	RA2 = RB*(1-2DC)/DC
Lower Resistor:	RB1 = DC*RA/(1-2*DC)
or	RB2 = TT*DC/(.693*TC)
or	RB3 = (TT-.693*RA*C)/(1.386*C)

The resistance equations to be used depend on data input and the sequence in which calculations are made. This is seen in the table below.

MENU SOLVE PROGRAM VARIABLE ASSIGNMENTS

SEL'NS FOR	n1 n2 n1 m2 m3	ms1 ms2	sel se2 se3
A,B TT,DC	0 1 2 3 4	0(TT) 1(DC)	TC RA RB
A,C			
A,D TT,RA	0 3 1 2 4	3(RA2) 0(TT)	DC TC RB
A,E TT,RB	0 4 1 2 3	4(RB1) 0(TT)	DC TC RA
B,C DC,TC	1 2 0 3 4	2(TC) 1(DC)	TT RA RB
B,D DC,RA	1 3 0 2 4	5(RA1) 1(DC)	TT TC RB
B,E DC,RB	1 4 0 2 3	7(RB3) 1(DC)	TT TC RA
C,D TC,RA	2 3 0 1 4	3(RB2) 2(TC)	TT TC RB
C,E TC,RB	2 4 0 1 3	4(RB1) 2(TC)	TT DC RA
D,E RA,RB	3 4 0 1 2	6(RB2) 3(RA2)	TT TC DC

Table 2. 555/556 timer IC equations and variable assignments as used in TIMR.C.

The below reprint is from National Semiconductor Linear Databook. It is application information typical of manufactures specifications used to calculate timer values.

LM555/LM555C

applications information

MONOSTABLE OPERATION

In this mode of operation, the timer functions as a one-shot (*Figure 1*). The external capacitor is initially held discharged by a transistor inside the timer. Upon application of a negative trigger pulse of less than 1/3 V_{CC} to pin 2, the flip-flop is set which both releases the short circuit across the capacitor and drives the output high.

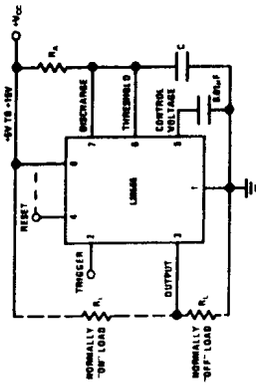

FIGURE 1. Monostable

The voltage across the capacitor then increases exponentially for a period of $t = 1.1\ R_A C$, at the end of which time the voltage equals 2/3 V_{CC}. The comparator then resets the flip-flop which in turn discharges the capacitor and drives the output to its low state. *Figure 2* shows the waveforms generated in this mode of operation. Since the charge and the threshold level of the comparator are both directly proportional to supply voltage, the timing internal is independent of supply.

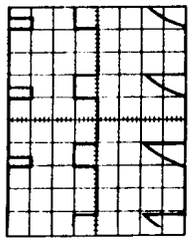

FIGURE 2. Monostable Waveforms

During the timing cycle when the output is high, the further application of a trigger pulse will not effect the circuit. However the circuit can be reset during this time by the application of a negative pulse to the reset terminal (pin 4). The output will then remain in the low state until a trigger pulse is again applied.

When the reset function is not in use, it is recommended that it be connected to V_{CC} to avoid any possibility of false triggering.

Figure 3 is a nomograph for easy determination of R, C values for various time delays.

ASTABLE OPERATION

If the circuit is connected as shown in *Figure 4* (pins 2 and 6 connected) it will trigger itself and free run as a

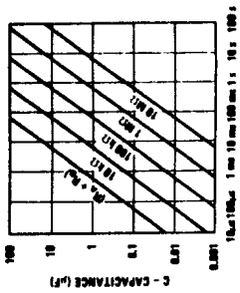

FIGURE 3. Time Delay

multivibrator. The external capacitor charges through $R_A + R_B$ and discharges through R_B. Thus the duty cycle may be precisely set by the ratio of these two resistors.

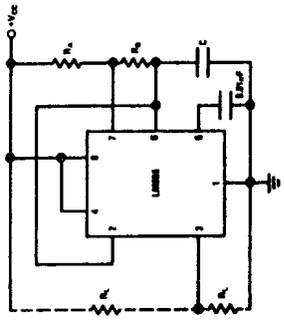

FIGURE 4. Astable

In this mode of operation, the capacitor charges and discharges between 1/3 V_{CC} and 2/3 V_{CC}. As in the triggered mode, the charge and discharge times, and therefore the frequency are independent of the supply voltage.

Figure 5 shows the waveforms generated in this mode of operation.

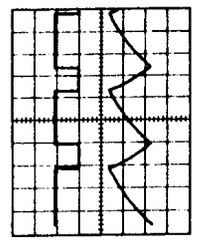

FIGURE 5. Astable Waveforms

The charge time (output high) is given by:

$$t_1 = 0.693\ (R_A + R_B)\ C$$

And the discharge time (output low) by:

$$t_2 = 0.693\ (R_B)\ C$$

Thus the total period is:

$$T = t_1 + t_2 = 0.693\ (R_A + 2R_B)\ C$$

Special Feature

Advanced Users

Forth Internals

FORTH MULTITASKING IN A NUTSHELL

By Brad Rodriguez

Multitasking in Forth is very simple, incredibly powerful...and widely misunderstood. This talk will describe what multitasking can do, and when it can be profitably used. The most common Forth multitasker -- the cooperative, round-robin model -- will be explored in general terms. Example code will be given for F83, which will be used to illustrate how to create and control parallel tasks, protect shared resources with semaphores, and pass messages between tasks. Preemptive and prioritized multitaskers will be mentioned briefly.

INTRODUCTION

What is multitasking?

Multitasking is the ability to run several independent programs in a single CPU, apparently simultaneously. Of course, a single CPU can only run one instruction at a time. The illusion is created by switching the CPU very quickly -- hundreds of times a second -- among several different programs. Each of these programs is called a "task," hence the name "multitasking."

A popular fallacy is that you need a big CPU, like a 68000, in order to do multitasking. This is simply not true. Any CPU can multitask; I've done it on Z80s and single-chip Z8s.

Another fallacy is that you need a real-time clock interrupt to "time slice" the tasks. While some multitaskers do indeed work this way, the Forth multitasker does not. You can multitask without interrupts, and with no special hardware support at all!

It is important to distinguish between multitasking and multiuser. Multiuser systems are those that support multiple terminals, and give several people the illusion of working each on his own computer. Multitasking is simpler: there is only one user, but he can have several things running simultaneously. Unix is multiuser. Amigas, Macintoshes, and Windows PCs are multitasking.

Occasionally, in a computer science text, you will see multitasking called multiprocessing. I prefer to avoid this usage. To me, multiprocessing means several CPUs operating in parallel. If it's on a single CPU, I call it multitasking.

When should multitasking be used?

There are several situations -- some more obvious than others -- which can benefit from multitasking.

Parallel operations. Some computer applications just naturally have several things running at once. An embedded process controller probably has to keep the control loops running, even while the operator is typing commands on the keypad. Network communications generally have to be maintained while other things are going on. And usually you'd like to be able to do something else while the printer is grinding out a long listing. All of these are candidates for multitasking.

Using idle time. Many computer programs spend a lot of time waiting -- maybe for a keypress, an external event, or a data transfer to complete. Sitting in a wait loop is a waste of good CPU time! Multitasking allows the CPU to be doing something else while waiting for an event.

Coroutines. Sometimes you need to break out of a program "in the middle." For example, I recently wrote a command processor which needed to stop and wait for a keypress at a deeply nested level. Multitasking allows a program to be suspended at any point, and later resumed, with no special effort on the part of the programmer.

Clarity and ease of programming. Often a program is easier to write, and its logic is more obvious, if you use the facilities that a multitasker gives you. Sometimes an algorithm is best expressed procedurally, in terms of waiting for an event -- even when you know the CPU can't just sit there and wait. I could have written the command processor using a huge state table, and entered the routine "at the top" on every keypress -- but it would have been much larger, and impossible to read and maintain. Multitasking adds **WAIT** and **PAUSE** to your "programming toolbox." You just use them naturally, and the multitasker takes care of the details.

HOW FORTH MULTITASKING WORKS "INSIDE"

What each task requires

What is needed to support multiple Forth tasks?

Separate programs. In most multitasking applications, each

task will run a different program. In single-user Forth systems, these need not be in separate dictionaries; they can simply be Forth words with different names.

In multiuser Forth systems, several people may be adding definitions at the same time. So, each user needs some RAM for a private dictionary. F83 does this. (How multiple dictionaries are managed, and linked to the "main" Forth dictionary, are beyond the scope of this article.)

Re-entrant kernel code. We'd really prefer not to have a separate copy of **DUP** or **FIND** for each task! Code which can be shared by several tasks simultaneously is called "re-entrant." Among other things, re-entrant code can't leave important data in global variables. Fortunately, Forth code, with its use of stacks, tends to be naturally re-entrant, and the "cooperative" multitasker relaxes some of the restrictions on temporary variables. Most Forth kernels are fully re-entrant.

Private stacks. Obviously, each independent Forth program will need its own parameter and return stacks. So, each task has some RAM for its stacks, and also has its own parameter and return Stack Pointers.

Private "user areas." Some things simply must be kept in variables, and yet will have different values for each task. The "bottom" stack addresses **S0** and **R0** are two examples; the **BASE** for number conversion is another. So, each task needs a private RAM area for variables, called the user area. Different Forth kernels may keep from six to several dozen variables in the user area. Most Forths provide a "user variable" which, instead of returning an absolute address, indexes into the user area of whichever task is currently running.

The user area is essential to hold certain task control data: among other things, the "saved" stack pointer, and the link to the next task. Both of these will be discussed shortly.

Private buffers. Certain buffers must exist privately for each task. One example is the **PAD** buffer which is used for numeric output. If two tasks tried to display a number at the same time, using a single **PAD**, nonsense would result! Such buffers may be kept in the user area, or (in multiuser systems) in the private dictionary.

Multiuser systems also need separate Terminal Input Buffers. Usually, though, a single set of disk buffers is shared by all the tasks.

Switching tasks

Most Forth systems use the simplest of multitasking schemes: the "round-robin, cooperative" multitasker. Round-robin means that each task takes its turn at the CPU, one at a time, in a fixed sequence -- a big loop of tasks. Cooperative means that each task has the CPU as long as it wants, and releases the CPU only when it's ready -- nothing will "grab" the CPU away from a task.

Figure 1 shows the RAM allocation for a three-task Forth system. The private areas described above are usually grouped together for convenience. The dictionary and the disk buffers are common to all of the tasks.

Switching from one task to another -- that is, from one program to another -- requires three steps.

a) Save the state of the current program. Everything necessary to restore this program -- to exactly the point where it was suspended -- must be saved. This is called the "task context," and may include the CPU's program counter, flags, and registers, as well as data in RAM.

In Forth, most of the task context is on the (private) parameter and return stacks, and so is safe from alteration. But there are four crucial values which are so frequently used that they are usually kept in CPU registers:

SP - the parameter Stack Pointer
RP - the Return stack Pointer
IP - the Interpreter Pointer
UP - the User Pointer (base address of the user area)

These four registers must be saved and restored when the task is switched. (It turns out that we don't need to save the CPU's Program Counter, since we never change tasks in the middle of a **CODE** word.)

b) Select the next task to run. This may be done in fixed rotation, or according to some priority scheme.

c) Restart the new task according to its saved context. This will resume execution of the new task at the point where it was last suspended.

The task switch in F83

The Forth word which switches tasks is traditionally called **PAUSE**. It must do the following:

a) Save the IP, RP, and SP. Typically two of these will be pushed on a stack, and that stack's pointer will then be saved in the user area. We don't need to save UP; you'll see why in a moment.

b) Get the address of the next task's user area. This is what the LINK variable in the user area is for: it contains the address of the next task in the round-robin sequence. (The tasks are thus chained together in a linked list.) Note that this link gives you the UP (user area address) for the new task! This is why UP doesn't need to be explicitly saved.

c) Restore the IP, RP, and SP of the new task.

d) Continue Forth execution!

The time it takes to do this is called the "context switch time," and is an important figure of merit for multitasking systems. Since Forth systems only have to save three CPU registers, their context switching times are quite fast -- under 10 microseconds on some 8-bit CPUs!

F83 on the IBM PC uses some tricks to improve performance. (Refer to Figure 1.) In F83, step (b) is performed by jumping into the next task's user area. In this area is a code fragment (INT 80h) which executes a **RESTART** routine. This **RESTART** routine does the step (c) described above.

Why does F83 do this? A task which is not ready to run can be "put to sleep" by changing the INT 80h to a JMP instruction. Then, when **PAUSE** jumps to this task, it immediately jumps to the next task. The sleeping task is skipped in only one machine instruction!

You can see that if all the tasks had JMP instructions, the round-robin loop would be simply a loop of jumps. Of course, the CPU would then be stuck in an infinite loop! Usually, several tasks will be "awake" with INT 80h instructions.

F83's **PAUSE** and **RESTART** are coded as follows:

```
CODE (PAUSE)
        IP PUSH             save IP on parameter stack
        RP PUSH             save RP on parameter stack
        UP #) BX MOV  SP 0 [BX] MOV  save SP in user area
        BX INC  BX INC  BX INC  BX INC  calculate address of next
        0 [BX] BX ADD  BX INC  BX INC    task from (relative) LINK,
        BX JMP  C;          then jump to that task

CODE RESTART                entered from an INT 80h instruction
        -4 # AX MOV
        BX POP              get return adrs saved by INT 80h
        AX BX ADD           adjust it to the start of user area
        BX UP #) MOV        make this the current UP
        AX POP  AX POP      clean up INT 80h leftovers
        STI
        0 [BX] SP MOV       restore SP from user area
        CLI
        RP POP              restore RP from parameter stack
        IP POP              restore IP from parameter stack
        NEXT  C;            continue Forth execution at new IP
```

USING A FORTH MULTITASKER

Creating a task

All Forth systems start with one task. (There's always at least one program running!) New tasks can then be added to the system in two phases, which I call creation and activation.

Creating a new task involves two steps:

a) Reserve RAM for the task. Space must be allocated for its two stacks and its user area. In multiuser systems, a private dictionary must be reserved as well.

b) Link the task into the round-robin list. This is a simple linked-list insertion using the LINK field.

These steps are performed only once. Obviously, reserving RAM twice for the same task is a pointless waste of RAM...and can lead to confusion if other tasks need to know where this task is located. Linking the task into the list twice is more subtle, and more insidious: usually it ends up destroying the round-robin loop!

In F83 on the IBM PC, steps (a) and (b) are performed by the word **TASK:**. This word expects, on the stack, the number of bytes to reserve for the new task. 256 bytes are reserved for the return stack, and the rest is divided between the user variables, the private dictionary, and the parameter stack. Thus:

```
HEX 400 TASK: SCREEN-CLOCK
```

defines a task and allocates a total of 1024 bytes to it. The newly defined word **SCREEN-CLOCK** will return the base address of this 1024-byte area (the base address of the user area).

Activating a task

Once the task has been created, it can be "activated" any number of times. This involves two steps:

c) Initialize the task context. Several user variables must be initialized, and in F83 the INT 80h instruction must be inserted in the user area. Also, the initial values for the SP and RP must be stored in the stacks or user area in such a way that they will be correctly loaded into the CPU registers when this task is "resumed" for the first time.

d) Specify the code to be executed by that task. The initial value for the IP must be stored in the stacks or user area, too. When the task is "resumed" for the first time, this will be where execution begins. It must point to a fragment of high-level Forth code.

(Some Forths may only perform step (c) once, in which case it may be done as part of task creation.)

After the task has been activated, it will lie dormant until its turn in the round-robin loop. Then it will begin executing the Forth code specified in step (d).

In F83, the RP, SP, and IP are initialized by the word **ACTIVATE**. F83's **ACTIVATE** must be used within a Forth word -- it cannot be used interpretively. It expects the address of the task on the stack, and is immediately followed by the high-level code the new task is to run. For example:

```
: START-CLOCK  SCREEN-CLOCK ACTIVATE   BEGIN .CLOCK PAUSE AGAIN ;
<-------------------------------------> <--------------------->
code performed by "main" task           code performed by
                                        SCREEN-CLOCK task
```

START-CLOCK is executed by some other Forth task -- typically the "main" (initial) task. It then "activates" the **SCREEN-CLOCK** task to perform the code fragment **BEGIN .CLOCK PAUSE AGAIN**. The main task exits this word immediately after **ACTIVATE**.

Other task control in F83

The functions described above -- **PAUSE**, **TASK:**, and **ACTIVATE** -- are all that you need to establish a multitasking Forth system. But F83 provides some additional words for convenience:

taskaddr SLEEP puts a task to sleep, by inserting the JMP instruction in the user area.

taskaddr WAKE awakens a task, by inserting the INT 80h instruction.

STOP just puts the running task to sleep, and then switches to the next task. This is equivalent to **my-task SLEEP PAUSE**.

MULTI enables the multitasker.

SINGLE disables the multitasker, by changing the action of **PAUSE** to a "no-op." (**PAUSE** is a DEFERred word for just this purpose.) Whichever task does **SINGLE** will keep control of the CPU.

APPLYING THE FORTH MULTITASKER

The PAUSE : running programs in parallel

The simplest use of a multitasker is to have several programs running in parallel. This requires only that every program have its own task, and that every program obey two rules:

a) Every task must **PAUSE** periodically! This is how the task voluntarily surrenders the CPU to the other tasks in the system. If there are no **PAUSE**s, this task will never release the CPU, and no other task will ever run!

Most Forth I/O words, such as **EMIT**, **KEY**, and **BLOCK**, contain a **PAUSE**. The assumption here is that I/O is slow, and so other tasks should have some time at the CPU.

b) The code executed by the new task must never return! Remember, this code was not entered from a subroutine call -- it is the very first code executed by this task. Thus, there is no return information on the return stack! (Boom.) In general, the "topmost" Forth code of any task must be an endless loop. In F83, the word **STOP** can be used to end a task, instead of looping.

There is also a rule for Forth programmers during the test and development phase: do not **FORGET** tasks! Remember, once a task is created, it is linked into the round-robin list. If you -- accidentally or intentionally -- reclaim a task's RAM area with **FORGET**, and then put other definitions into that RAM, you will destroy the round-robin linked list! (Again, boom.) It's safest to put all the task creation first in your code, and then never **FORGET** back that far.

Example #1: the on-screen clock

Note: the example code given in this article is written for F83 version 2.1.0 for the IBM PC. F83.COM seems to be distributed with the multitasker already installed. Type

' TASK: .

to find out. If **TASK:** is not defined, you will need to install the multitasker by typing

OPEN CPU8086.BLK 21 LOAD OPEN UTILITY.BLK 52 LOAD

Screen 1 of the listing is the code for a simple on-screen clock. This code creates a second task which continually displays the time of day in the upper right-hand corner of the screen.

@TIME is a Forth word to return the IBM PC clock time.

We want to use the Forth display words, but we don't want the clock task to interfere with the main task's display. After we reposition the cursor to the top right of the screen, we need to be able to put it back where it was. The words **@CURS** and **!CURS** invoke BIOS functions to get and set the current display cursor.

.TIME displays the time in hh:mm:ss format. It illustrates the use of **@CURS** and **!CURS** to get the cursor position, set a new position, and restore the original position. Note the use of **SINGLE** and **MULTI** around **TYPE**. **TYPE** does many **EMIT**s, and each **EMIT** does a **PAUSE**. This would switch back to the main task while the cursor is in the wrong position! Rather than redefine **EMIT** to eliminate the **PAUSE**, we can simply shut off multitasking for the duration of the **TYPE**. (This is, however, quite crude.)

The definition and activation of the **SCREEN-CLOCK** task have already been described. Note that you must specify **DECIMAL** from within the clock task's code. The number base is a user variable, and typing **DECIMAL** from the keyboard will change the main task's number base, not the clock task's.

After loading this screen, type

START-CLOCK MULTI

to activate the clock. You will see the clock appear in the corner of the screen. Type **WORDS** and observe that the displays do not interfere with each other. Then type **SINGLE WORDS** and compare the speed when the multitasker is

switched off. The clock display in this example consumes far too much CPU time; it redisplays the clock dozens of times every second, when only once per second would be adequate. A better version would wait for the time to change before redisplaying.

Example #2: the round-robin cycle counter

Screen 2 is Forth code to count passes through the round-robin list. Similar code is supplied as an example with F83; this code is different only in that it displays the cycle count in the upper right corner. Load this screen, and type

START-COUNTER

to activate the counter. (You may have to turn **MULTI** back on.) On my 12 MHz AT, with the on-screen clock task also running, I see about 50 counts per second. This means that each task is getting the CPU every 20 milliseconds.

Waiting without pain

Programmers seem to write a lot of wait loops. Most wait loops fall into one of three categories:

a) polling an I/O device. Sometimes you must use polled I/O. The hardware may not have a "data ready" interrupt. Operating system software may only offer a "check status" function -- such as the keyboard under MS-DOS.

b) awaiting an interrupt. When the hardware supports interrupts, you may have to wait for an interrupt to occur. For example, a disk controller using DMA will issue an "end of transfer" interrupt when the operation is complete.

c) waiting for another task. In multitasking applications, occasionally one task has to wait until another task has accomplished something.

Sometimes, waiting is just the obvious way to write a program! Recall the example of the command processor, which has to wait for a keypress to be received and processed.

Wait loops burn up CPU time. Even worse, trying to wait for more than one thing at a time can lead to some Byzantine compound loops! A multitasker solves both of these problems.

One simple addition changes the wait loop from a CPU hog to an efficient programming construct. Simply insert a **PAUSE** in the loop! This ensures that, while this task is waiting, all the other tasks in the system will get to use some CPU time.

Most multitasking Forth kernels write their I/O this way. For example, the basic definition of **KEY**

```
: (KEY)  BEGIN (KEY?) UNTIL  0 8 BDOS ;
```

is changed in F83 to

```
: (KEY)  BEGIN PAUSE (KEY?) UNTIL  0 8 BDOS ;
```

As long as no keypress is ready, this task will **PAUSE** repeatedly, surrendering the CPU to the other tasks in the system. The polling represents a slight overhead: on every pass through the task list -- typically every few milliseconds -- this task will execute **PAUSE (KEY?) UNTIL**. If **(KEY?)** is not too complex, this overhead is negligible.

Note that, even if a keypress is ready, the word **(KEY)** will do at least one **PAUSE**. This is usually desirable. If this is not desirable, a **BEGIN..WHILE..REPEAT** loop can be used instead.

What if all of the tasks are waiting for something? Then the system ends up spinning in a much larger wait loop -- checking each task in turn for its "wake up" condition, and moving on to the next. The first task whose poll is successful will then continue execution. With no special programming effort, the multitasker automatically checks a set of events, and starts a different program for each event!

Be aware that putting a **PAUSE** in the loop will significantly reduce the polling rate -- from microseconds to milliseconds. If fast response is essential, or there is a chance that an event can be "missed," some other approach is needed. Interrupts are usually best in this case.

Shared resources and semaphores

The problem of mutual exclusion

Once you have several tasks running in parallel, a new problem can arise: access conflicts. This happens when two tasks attempt to use, simultaneously, some device or resource which can only be used by one task at a time.

Consider, for instance, a disk controller which requires two time-consuming operations to access the disk: a seek, and then a read or write. (Many controllers are like this; the Western Digital 1791 is one example.) Suppose Task A does a seek, which takes a few hundred milliseconds. While it is waiting, other tasks are running. Now suppose that Task B tries at this moment to access the disk. It issues its own seek command, conflicting with the command issued by Task A. When the seek completes, and Task A resumes, it will be at the wrong location on the disk. (Worse, since only one completion signal is issued by the controller, one of the tasks may wait forever for its seek to finish.)

Or consider the on-screen clock and cycle counter examples: while one task had altered the display cursor, other tasks had to be prevented from using it.

Or consider the printer. If one task -- maybe a background printing task -- is outputting to the printer, we certainly don't

want another task sending output at the same time.

In the example programs we used the crude solution of switching off the multitasker. This is usually not practical -- it defeats the whole purpose of multitasking! We need a better solution.

Fortunately, this "mutual exclusion" is a classic problem in computer science, and several solutions have been devised over the decades [TAN87]. One of the simplest and most elegant is the "semaphore," invented by by E. W. Dijkstra in 1965.

The semaphore

In its simplest form, a "binary semaphore" is a flag associated with a resource. Two operations act on semaphores: **WAIT** and **SIGNAL**. **WAIT** checks to see if the resource is available. If so, it is marked "unavailable"; if not, the CPU is released to other tasks until the resource becomes available. **SIGNAL** just marks the resource "available."

In Forth, these can be written

```
: WAIT ( addr -- )
        BEGIN PAUSE DUP C@ UNTIL \ wait for nonzero = available
        0 SWAP ! ;                \ make it zero = unavailable

: SIGNAL ( addr -- )
        1 SWAP ! ;          \ make it nonzero = available
```

(These are also in screen 3 of the listing.)

Dijkstra observed that the operations of testing and setting the semaphore must be indivisible. In the cooperative Forth multitasker, all Forth code is indivisible until a **PAUSE** is executed, so these definitions will work as written. In preemptive multitaskers, or any application where interrupts can affect semaphores, interrupts must be disabled within **WAIT** and **SIGNAL** (except for the **PAUSE**).

It is a subject of some debate whether **SIGNAL** should include a **PAUSE**. Adding a **PAUSE** ensures that a task which is blocked on that semaphore will be awakened soonest, thus maximizing the use of the resource. Sometimes, however, this is not what is desired.

How semaphores are used

Every task, before using a shared resource, does a **WAIT** on its semaphore, and after using the resource, does a **SIGNAL**. This is sufficient to ensure that only one task can use that resource at any time -- and yet even if one task is blocked, the other tasks can run normally.

The semaphore is defined as a simple Forth variable. The semaphore variable must be initialized to a nonzero value, to indicate that the resource is initially "available." This should be done when all other variables are initialized: when you load the application (on a PC), or as part of the startup code (in an embedded application.)

For the disk example described above, the Forth code may look something like

```
VARIABLE DISK-SEMAPHORE

: READ-SECTOR
        DISK-SEMAPHORE WAIT
        SEEK
        READ
        DISK-SEMAPHORE SIGNAL ;
```

Now consider the situation with two tasks trying to access the disk simultaneously:

```
Task A                  Task B

WAIT

SEEK (100 msec)
                        WAIT
READ (100 msec)         |
                        |---task #2 is blocked!
SIGNAL                  |
                        SEEK (100 msec)

                        READ (100 msec)

                        SIGNAL
```

Task A performs a **WAIT**, sees that the disk is available, and proceeds to do its **SEEK**. Task A then waits with **PAUSE** until the seek is complete, so that other tasks can use the CPU. At this point Task B requests the disk resource with **WAIT**. Thanks to Task A, the resource is busy, so Task B is put into a wait loop -- waiting not for the seek to complete, but for **DISK-SEMAPHORE** to be set. This won't happen until Task A is finished seeking and reading. So, Task A is waiting for the disk, Task B is waiting for Task A's **SIGNAL**, and the other tasks in the system are free to use the CPU.

The beauty of this approach is that the **WAIT** and **SIGNAL** are identical for all tasks, so they can be made part of the common routine **READ-SECTOR** that all tasks use.

Semaphores for intertask signalling

The **WAIT** operator is general-purpose: it can be used any time you need to wait for a RAM location to become nonzero. Recall that two of the uses of wait loops are to wait for an interrupt to occur, and to wait for a signal from another task.

To wait for an interrupt, simply **WAIT** on a variable you have defined for the purpose. The interrupt routine then just needs to store a nonzero value in that variable. (Often this can be done with a single machine instruction, e.g., increment.) Waiting for another task is exactly the same, except that the other task can use the high-level **SIGNAL** word to store the non-zero value.

Note that, in these cases, the semaphore should be initialized to a zero value.

Problems with semaphores

Semaphores are not foolproof. They are susceptible to "deadlock," where two tasks, each having grabbed one of two resources, are waiting for the other resource to become available. One solution to this is the "monitor," also described in [TAN87], which can be implemented if you have semaphores.

Breaking out of the depths

What about the problem of "breaking out" of a deeply nested program, preserving all of its nesting information, and then returning to it at some future time?

Surprise! This is exactly what **PAUSE** does! All of the return information, temporary variables, etc. which need to be preserved, are all part of the task context. The task context, you will recall, is the information which is saved when tasks are switched. So, to suspend a task until some future event, simply do a **WAIT** or some other multitasking wait loop.

Returning to the command processor example: when I rewrote this application, I created a word **KEYPRESS** which waited -- with **PAUSE**s -- for a keypress to become available. I also converted all of the variables used in command processing to user variables. **KEYPRESS** is used in dozens of places in the program, as the command processor works its way through its logic tree. The command processing code is written "normally," with no special attention given to saving the state, or releasing the CPU to other tasks!

Message passing

Often parallel tasks need to communicate with each other. One of the most-used schemes is the "mailbox," an agreed-upon place where messages can be left for a task. Mailboxes may be statically allocated to each task -- in which case you must be careful to have enough mailboxes -- or they can be dynamically allocated from a pool.

Perhaps the simplest scheme, which ensures that there are sufficient mailboxes for all the tasks, is to include a mailbox in each task's user area. Then, as tasks are added, new mailboxes are too. Each task can reach its own mailbox easily, as a user variable. Other tasks can reach that mailbox by offset from the given task's address. (Recall that the "task address" is usually the base address of its user area.)

For simplicity, we can limit messages to a single cell (16 bits on the IBM PC). An F83 implementation of this would be:

```
: SEND  ( message taskadr -- )
        MYTASK
        OVER SENDER LOCAL  get adrs of destination's SENDER
        BEGIN
                PAUSE              loop with PAUSE,
                DUP @              until his SENDER is zero
        0= UNTIL
        !                          store my adrs in his SENDER
        MESSAGE LOCAL ! ; store the message in his MESSAGE

: RECEIVE  ( -- message taskadr )
        BEGIN
                PAUSE              loop with PAUSE,
                SENDER @           until my SENDER is nonzero
        UNTIL
        MESSAGE @       get the message from my MESSAGE
        SENDER @        get his task adr from my SENDER
        0 SENDER ! ;    indicate mailbox empty & ready
```

MESSAGE and **SENDER** are user variables. (See screen 4 of the listing for their definition.) The convention here is that, if **SENDER** is nonzero, a message is in the task's mailbox. This works because most Forth systems can't have a task user area at address 0000.

SEND sends the given message to the given task. Before it can do so, it must be sure that the destination mailbox is empty -- otherwise it would destroy some other mail. So it waits, with **PAUSE**s, until the destination task's **SENDER** variable is zero. Then it stores the message, and the sending task's address (**MYTASK**), in the destination user area.

RECEIVE waits for any message to appear, as indicated by **SENDER** going nonzero. It then reads the message and the sending task's address. Finally, it clears the **SENDER** field, to indicate that the mailbox is now empty.

Note that if two tasks try to send messages to the same destination, one of the tasks will be blocked until the destination task has read its mail!

This is a very simplistic message system, although adequate for many applications. More sophisticated schemes allow a receiver to select only messages from a specific sender, or allow multiple messages to be queued at the receiver so that multiple senders do not have to wait. Other extensions would allow variable-length or string messages.

Example #3: the Print utility

Screen 5 shows a simple use of messages. This is a modification of F83's screen-printing utility **SHOW**, to run as a "background" task. Load this screen, and type

START-SHOW

to activate the printing task. (Don't forget to turn **MULTI** on.) At this point, nothing will happen. You can now type

1 4 SHOW

and a printout of screens 1 to 4 will start on the printer. At the same time, you will get the "ok" prompt on the screen, and you can type Forth commands. Type **WORDS**, for instance.

When the printing task was activated, its first action (in **START-SHOW**) was to clear its mailbox. It then waits for two messages in a row, which it expects to be the first and last screen to print, respectively. (We don't care who sends these messages, so the sending task address is **DROP**ped after each

RECEIVE.) Since no message has yet been sent, it just waits forever, invisible to the "main" Forth task.

SHOW is redefined to simply send two messages to the printing task. So, when the **SHOW** command is typed at the keyboard, the printing task is awakened and begins to print the requested screens. When the listing is complete, the printing task loops, waiting for a new print request (two more messages).

If you have a slow printer with a small buffer, you will notice occasional interruptions in the **WORDS** listing being displayed by the main task. This might appear as if the multitasker is not working as advertised. Actually, this is an unfortunate byproduct of the MS-DOS printer routine. MS-DOS assumes that a program which is printing can afford to wait, so there is no BIOS call to test the "printer ready" flag. All Forth can do is request a character be output to the printer. If the printer is not ready, MS-DOS (not Forth) will tie up the computer while it waits for the printer. To fix this problem, we would have to bypass MS-DOS and write our own direct printer I/O routines.

Advanced topics

The round-robin, cooperative multitasker which has been described is the most common Forth multitasker. Most public-domain Forths and many commercial Forths use this approach -- it's simple, versatile, and efficient. But some Forth systems have extended the multitasker to add new capabilities; and most non-Forth multitaskers offer some of these "advanced features."

Preemptive multitaskers

The problem with cooperative multitaskers is that each task keeps the CPU until it does a **PAUSE**. It can be very difficult to distribute CPU time equally among the tasks; poorly written code can "hog" the CPU for inordinate amounts of time. Placement of **PAUSE** instructions becomes something of an art.

One solution is to allow a task switch to be forced upon a task, say by a real-time clock interrupt. If tasks are switched every 10 milliseconds, it's easy to allocate CPU time accurately. This also ensures that no task can tie up the CPU; and in a round-robin system, it guarantees an upper bound on how long a task must wait for the CPU.

Preemptive multitaskers have several disadvantages, too:

a) since a task switch can occur at any time -- even in the middle of a **CODE** word -- all of the CPU registers and flags need to be saved with the task context. On a big processor like the 68000, this can make the context switch time much longer. (Recall that the cooperative multitasker only needs to save four registers.)

b) words such as **WAIT** and **SIGNAL** become more complicated, since they must switch interrupts off to be indivisible.

c) since Forth code is no longer "normally" indivisible, many more semaphores are needed to enforce mutual exclusion, and more temporary variables must be moved into the user area.

Prioritized multitaskers

Another problem is: how do you tell the system that some tasks are more important than others?

Perhaps you want to allocate more CPU to one task than another. Or, it may be desirable to run one task in preference to all others, as long as it's not waiting for an event. Or possibly there is a "background" task which should only run when no other task is able to run (i.e., when all the other tasks are waiting for something). What you need is a way to assign priorities to tasks. High-priority tasks should run in preference to low-priority tasks -- either by receiving a bigger share of the CPU time, or by taking the CPU completely away from lower-priority tasks.

Priority schemes are usually associated with preemptive multitaskers, but this need not be the case. Here's one possible modification to the basic Forth multitasker to implement a priority scheme: instead of **PAUSE** switching to the next task in the round-robin list, make it switch to a designated "first" task in the list. This task then is the highest-priority task; as long as it is able to run, no other task will get service. **WAIT** must be modified to link to the "next" task in the round-robin list.

The disadvantages of priority schemes are:

a) there's usually more overhead involved in task switching, since the multitasker must decide which task to run next.

b) assigning priorities is not simple. Sometimes a small change can have a disproportionate effect on the amount of CPU time various tasks receive.

c) in some systems, it's possible that some tasks will never get any CPU time.

Interrupts

How interrupts are handled can vary from one multitasker to another.

The simplest approach is to treat interrupts completely apart from the multitasker. Interrupts do not affect the execution of tasks; when an interrupt occurs, the interrupt routine executes, and then returns to whichever task was running. (I.e., interrupts work the same as they do in a non-multitasking system.) Interrupts can set flags which will later be seen by various tasks. This approach works best when the interrupt handlers can be kept simple. It offers the fastest interrupt service time;

however, if the purpose of the interrupt is to "wake up" a task, it may be many milliseconds before that happens.

A preemptive multitasker can allow an interrupt to wake up a specific task. If the service to be performed is complex -- say, interpreting a long message received via DMA -- then it's usually better to do that service in a task, rather than in an interrupt handler. (This is because, usually, some or all interrupts are blocked for the duration of an interrupt handler.) The downside: in addition to the disadvantages of preemptive taskers, this approach has a longer interrupt service time, since the interrupt handler is more complex.

There are other variations, but most interrupt service schemes fall loosely into one of these two categories.

Conclusion

Multitasking is a powerful tool, and like a good tool, it can dramatically increase your productivity. The basic Forth multitasker is more than adequate for many applications. Like Forth itself, it is simple, versatile, and efficient...and easy for one person to grasp. Add multitasking to your "toolkit": you will find that many programming problems become simpler, and some become trivial!

References

[LAX84] Laxen, H. and Perry, M., F83 for the IBM PC, version 2.1.0 [1984]. Distributed by the authors.

[TAN87] Tanenbaum, Andrew S., Operating Systems: Design and Implementation, Prentice-Hall [1987], 719 pp. Any worthwhile book on operating systems should describe semaphores. This book is recent, quite complete, and very readable. It has an good discussion of the various problems of mutual exclusion.

[TIN86] Ting, C.H., Inside F83, 2nd ed., Offete Enterprises, 1306 South "B" Street, San Mateo, CA, 94402, USA [1986], 287 pp. This explains the F83 multitasker quite well. Actually, it explains most of F83 quite well: a "must" for F83 users.

F83 for the IBM PC, CP/M, and 68000 are available on the Forth Roundtable on GEnie. F83 for the IBM PC is also available from the Forth Interest Group, P.O. Box 2154, Oakland, CA 94621, and from many shareware vendors.

The source code for this article is available on GEnie as file MULTIDEM.BLK.

T-Recursive Technology
221 King St. East, #32
Hamilton, ON L8N 1B5 Canada


```
Screen 1
========
\ MULTITASKER EXAMPLE 1 - ON SCREEN CLOCK          bjr09sep92
CODE @TIME ( -- h m s)  HEX 2C # AH MOV  21 INT  AH AH SUB
  CH AL MOV  AX PUSH  ( hrs)  CL AL MOV  AX PUSH  ( mins)
  DH AL MOV  1PUSH  ( secs)  END-CODE
CODE @CURS ( -- n)  0 # BH MOV  3 # AH MOV  10 INT  DX PUSH
  NEXT  END-CODE
CODE !CURS ( n --)  DX POP  0 # BH MOV  2 # AH MOV  10 INT
  NEXT  END-CODE
: .TIME ( h m s --)  0 <# # # 3A HOLD  NIP # # 3A HOLD  NIP
  # # #>  @CURS >R  48 !CURS  SINGLE TYPE MULTI  R> !CURS ;

HEX 400 TASK: SCREEN-CLOCK
: START-CLOCK  SCREEN-CLOCK ACTIVATE
  DECIMAL BEGIN @TIME .TIME PAUSE AGAIN ;
: STOP-CLOCK  SCREEN-CLOCK SLEEP ;
DECIMAL

Screen 2
========
\ MULTITASKER EXAMPLE 2 - ROUND-ROBIN CYCLE COUNTER

VARIABLE CYCLES  0 CYCLES !  HEX
: .CYCLES  1 CYCLES +!  CYCLES @
  @CURS >R  40 !CURS  SINGLE 6 U.R MULTI  R> !CURS ;

HEX 400 TASK: CYCLE-COUNTER
: START-COUNTER  CYCLE-COUNTER ACTIVATE
  DECIMAL BEGIN .CYCLES PAUSE AGAIN ;
: STOP-COUNTER  CYCLE-COUNTER SLEEP ;
DECIMAL

Screen 3
========
\ SEMAPHORES IN F83                          bjr09sep92
: WAIT ( addr -- )       \ WAIT for semaphore ready
  BEGIN PAUSE DUP C@ UNTIL  \ wait for nonzero = available
  0 SWAP ! ;             \ make it zero = unavailable

: SIGNAL ( addr -- )     \ SIGNAL that semaphore is ready
  1 SWAP ! ;             \ make it nonzero = available

Screen 4
========
\ MESSAGES IN F83                            bjr09sep92
USER VARIABLE MESSAGE    \ a 16-bit message
   VARIABLE SENDER       \ holds address of the sending task
FORTH
: MYTASK ( -- a)  UP @ ;    \ returns addr of the running task

: SEND ( msg taskadr -- )    \ send msg to the given task
  MYTASK OVER SENDER LOCAL  \ -- msg taskadr mytask SENDERadr
  BEGIN PAUSE DUP @ 0= UNTIL \ wait until his SENDER is zero
  ! MESSAGE LOCAL ! ;     \ store mytask,msg in his user var

: RECEIVE ( -- msg taskadr)    \ wait for a message from anyone
  BEGIN PAUSE SENDER @ UNTIL  \ wait until my SENDER nonzero
  MESSAGE @  SENDER @        \ get message and sending task
  0 SENDER ! ;               \ now ready for another message!

Screen 5
========
\ MULTITASKER EXAMPLE 3 - BACKGROUND PRINT UTILITY

HEX 400 TASK: SHOW-TASK

: START-SHOW  SHOW-TASK ACTIVATE
  DECIMAL  0 SENDER !  \ must clear message buffer first
  BEGIN
    RECEIVE DROP  RECEIVE DROP  \ -- first last
    SHOW
  AGAIN ;

: STOP-SHOW  SHOW-TASK SLEEP ;

: SHOW ( first last -- )
  SWAP  SHOW-TASK SEND  SHOW-TASK SEND ;
DECIMAL
```

FIGURE 1. TASK CONTEXT

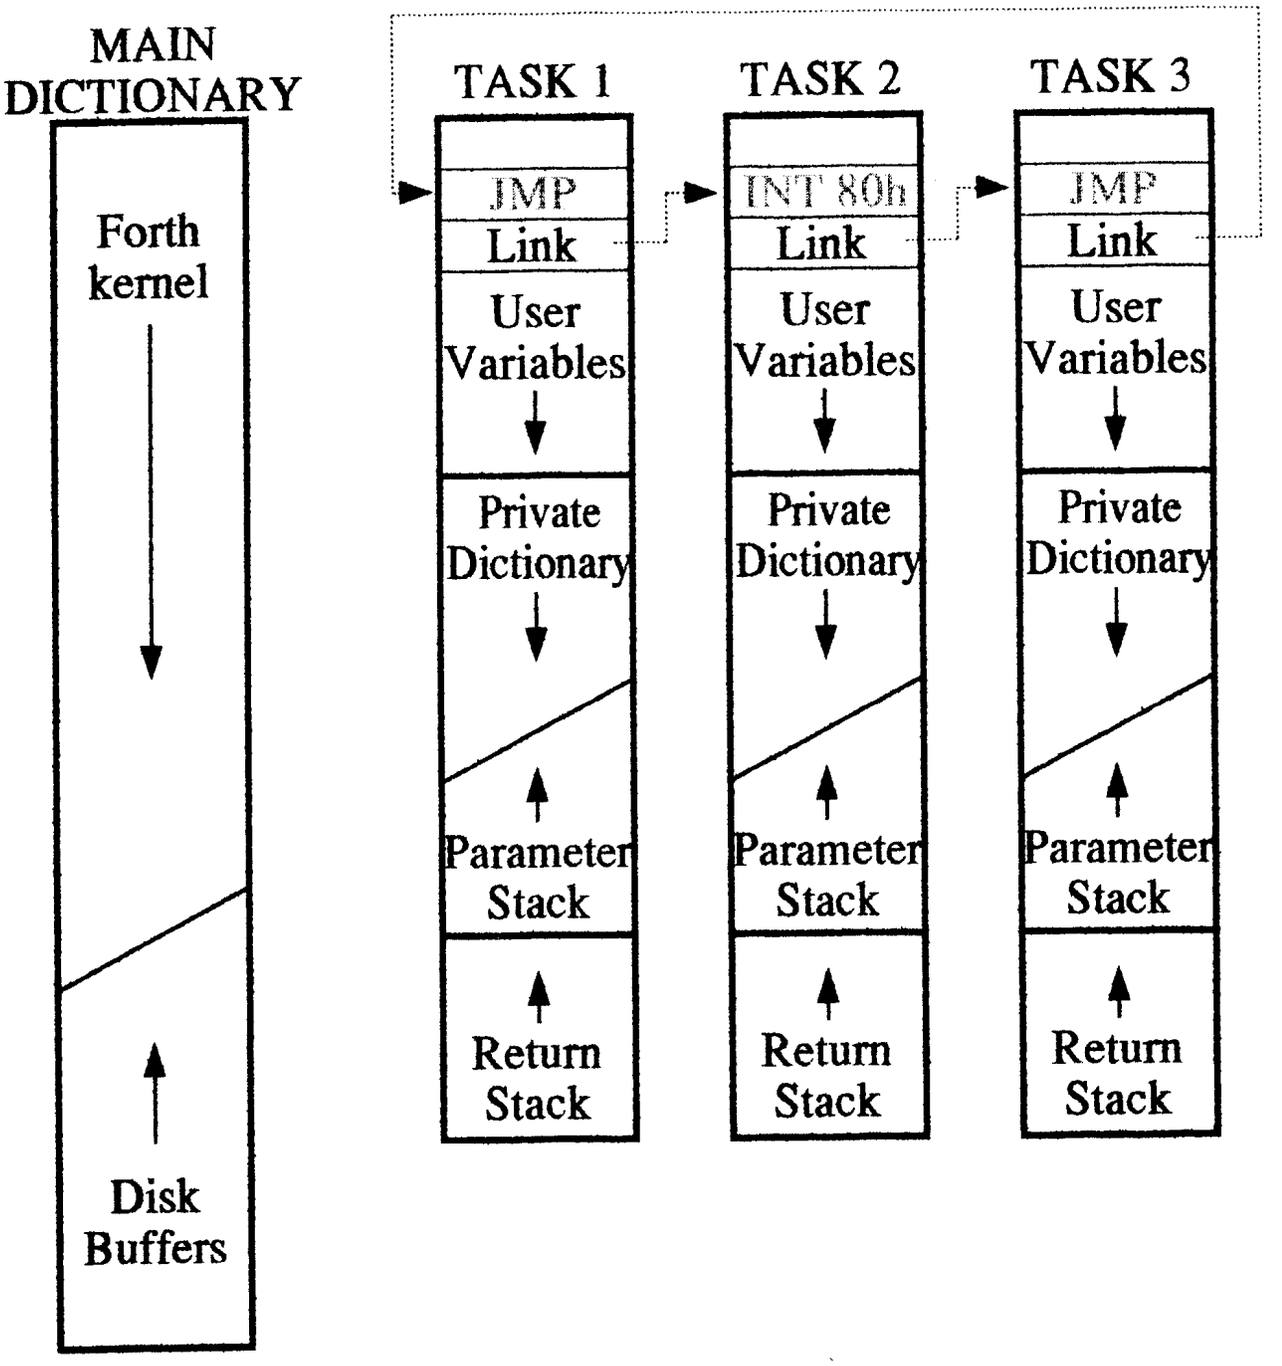

CPU REGISTERS

UP - User Ptr
SP - Stack Ptr
RP - Rtn Stk Ptr
IP - Interpret Ptr

The Computer Journal

Back Issues

Sales limited to supplies in stock.

Issues 1 to 19 are currently OUT of print. To assist those who want a full collection of TCJ issues we are preparing photo-copied sets. The sets will be issue 1 to 9 and 10 to 19. Each set will be bound with a plastic protective cover.

The price is $25 (in US) and $35 (foriegn air mail) . Expect TWO to THREE weeks for delivery after payment received at TCJ. Some single copies available, contact TCJ before ordering.

Issue Number 20:

- Designing an 8035 SBC
- Using Apple Graphics from CP/M: Turbo Pascal Controls Apple Graphics
- Soldering & Other Strange Tales
- Build an S-100 Floppy Disk Controller: WD2797 Controller for CP/M 68K

Issue Number 21:

- Extending Turbo Pascal: Customize with Procedures & Functions
- Unsoldering: The Arcane Art
- Analog Data Acquisition & Control: Connecting Your Computer to the Real World
- Programming the 8035 SBC

Issue Number 22:

- NEW-DOS: Write Your Own Operating System
- Variability in the BDS C Standard Library
- The SCSI Interface: Introductory Column
- Using Turbo Pascal ISAM Files
- The Ampro Little Board Column

Issue Number 23:

- C Column: Flow Control & Program Structure
- The Z Column: Getting Started with Directories & User Areas
- The SCSI Interface: Introduction to SCSI
- NEW-DOS: The Console Command Processor
- Editing the CP/M Operating System
- INDEXER: Turbo Pascal Program to Create an Index
- The Ampro Little Board Column

Issue Number 24:

- Selecting & Building a System
- The SCSI Interface: SCSI Command Protocol
- Introduction to Assemble Code for CP/M
- The C Column: Software Text Filters
- Ampro 186 Column: Installing MS-DOS Software
- The Z-Column
- NEW-DOS: The CCP Internal Commands
- ZTime-1: A Real Time Clock for the Ampro Z-80 Little Board

Issue Number 26:

- Bus Systems: Selecting a System Bus
- Using the SB180 Real Time Clock
- The SCSI Interface: Software for the SCSI Adapter
- Inside Ampro Computers
- NEW-DOS: The CCP Commands (continued)
- ZSIG Corner
- Affordable C Compilers
- Concurrent Multitasking: A Review of DoubleDOS

Issue Number 27:

- 68000 TinyGiant: Hawthorne's Low Cost 16-bit SBC and Operating System
- The Art of Source Code Generation: Disassembling Z-80 Software
- Feedback Control System Analysis: Using Root Locus Analysis & Feedback Loop Compensation
- The C Column: A Graphics Primitive Package
- The Hitachi HD64180: New Life for 8-bit Systems
- ZSIG Corner: Command Line Generators and Aliases
- A Tutor Program in Forth: Writing a Forth Tutor in Forth
- Disk Parameters: Modifying the CP/M Disk Parameter Block for Foreign Disk Formats

Issue Number 28:

- Starting Your Own BBS
- Build an A/D Converter for the Ampro Little Board
- HD64180: Setting the Wait States & RAM Refresh using PRT & DMA
- Using SCSI for Real Time Control
- Open Letter to STD Bus Manufacturers
- Patching Turbo Pascal
- Choosing a Language for Machine Control

Issue Number 29:

- Better Software Filter Design
- MDISK: Adding a 1 Meg RAM Disk to Ampro Little Board, Part 1
- Using the Hitachi hd64180: Embedded Processor Design
- 68000: Why use a new OS and the 68000?
- Detecting the 8087 Math Chip
- Floppy Disk Track Structure
- The ZCPR3 Corner

Issue Number 30:

- Double Density Floppy Controller
- ZCPR3 IOP for the Ampro Little Board
- 3200 Hackers' Language
- MDISK: Adding a 1 Meg RAM Disk to Ampro Little Board, Part 2
- Non-Preemptive Multitasking
- Software Timers for the 68000
- Lilliput Z-Node
- The ZCPR3 Corner
- The CP/M Corner

Issue Number 31:

- Using SCSI for Generalized I/O
- Communicating with Floppy Disks: Disk Parameters & their variations
- XBIOS: A Replacement BIOS for the SB180
- K-OS ONE and the SAGE: Demystifying Operating Systems
- Remote: Designing a Remote System Program
- The ZCPR3 Corner: ARUNZ Documentation

Issue Number 32:

- Language Development: Automatic Generation of Parsers for Interactive Systems
- Designing Operating Systems: A ROM based OS for the Z81
- Advanced CP/M: Boosting Performance
- Systematic Elimination of MS-DOS Files: Part 1, Deleting Root Directories & an In-Depth Look at the FCB
- WordStar 4.0 on Generic MS-DOS Systems: Patching for ASCII Terminal Based Systems
- K-OS ONE and the SAGE: System Layout and Hardware Configuration
- The ZCPR3 Corner: NZCOM and ZCPR34

Issue Number 33:

- Data File Conversion: Writing a Filter to Convert Foreign File Formats
- Advanced CP/M: ZCPR3PLUS & How to Write Self Relocating Code
- DataBase: The First in a Series on Data Bases and Information Processing
- SCSI for the S-100 Bus: Another Example of SCSI's Versatility
- A Mouse on any Hardware: Implementing the Mouse on a Z80 System
- Systematic Elimination of MS-DOS Files: Part 2, Subdirectories & Extended DOS Services
- ZCPR3 Corner: ARUNZ Shells & Patching WordStar 4.0

Issue Number 34:

- Developing a File Encryption System.
- Database: A continuation of the data base primer series.
- A Simple Multitasking Executive: Designing an embedded controller multitasking executive.
- ZCPR3: Relocatable code, PRL files, ZCPR34, and Type 4 programs.
- New Microcontrollers Have Smarts: Chips with BASIC or Forth in ROM are easy to program.
- Advanced CP/M: Operating system extensions to BDOS and BIOS, RSXs for CP/M 2.2.
- Macintosh Data File Conversion in Turbo Pascal.
- The Computer Corner

Issue Number 35:

- All This & Modula-2: A Pascal-like alternative with scope and parameter passing.
- A Short Course in Source Code Generation: Disassembling 8088 software to produce modifiable assem. source code.
- Real Computing: The NS32032.
- S-100: EPROM Burner project for S-100 hardware hackers.
- Advanced CP/M: An up-to-date DOS, plus details on file structure and formats.
- REL-Style Assembly Language for CP/M and Z-System. Part 1: Selecting your assembler, linker and debugger.
- The Computer Corner

Issue Number 36:

- Information Engineering: Introduction.
- Modula-2: A list of reference books.
- Temperature Measurement & Control: Agricultural computer application.
- ZCPR3 Corner: Z-Nodes, Z-Plan, Amstrand computer, and ZFILE.
- Real Computing: NS32032 hardware for experimenter, CPUs in series, software options.
- SPRINT: A review.
- REL-Style Assembly Language for CP/M & ZSystems, part 2.
- Advanced CP/M: Environmental programming.
- The Computer Corner.

Issue Number 37:

- C Pointers, Arrays & Structures Made Easier: Part 1, Pointers.
- ZCPR3 Corner: Z-Nodes, patching for NZCOM, ZFILER.
- Information Engineering: Basic Concepts: fields, field definition, client worksheets.
- Shells: Using ZCPR3 named shell variables to store date variables.
- Resident Programs: A detailed look at TSRs & how they can lead to chaos.
- Advanced CP/M: Raw and cooked console I/O.
- Real Computing: The NS 32000.
- ZSDOS: Anatomy of an Operating System: Part 1.
- The Computer Corner.

Issue Number 38:

- C Math: Handling Dollars and Cents With C.
- Advanced CP/M: Batch Processing and a New ZEX.
- C Pointers, Arrays & Structures Made Easier: Part 2, Arrays.
- The Z-System Corner: Shells and ZEX, new Z-Node Central, system security under Z-Systems.
- Information Engineering: The portable Information Age.
- Computer Aided Publishing: Introduction to publishing and Desk Top Publishing.
- Shells: ZEX and hard disk backups.
- Real Computing: The National Semiconductor NS320XX.
- ZSDOS: Anatomy of an Operating System, Part 2.

Issue Number 39:

- Programming for Performance: Assembly Language techniques.
- Computer Aided Publishing: The Hewlett Packard LaserJet.
- The Z-System Corner: System enhancements with NZCOM.
- Generating LaserJet Fonts: A review of Digi-Fonts.
- Advanced CP/M: Making old programs Z-System aware.
- C Pointers, Arrays & Structures Made Easier: Part 3: Structures.
- Shells: Using ARUNZ alias with ZCAL.
- Real Computing: The National Semiconductor NS320XX.
- The Computer Corner.

Issue Number 40:

- Programming the LaserJet: Using the escape codes.
- Beginning Forth Column: Introduction.
- Advanced Forth Column: Variant Records and Modules.
- LINKPRL: Generating the bit maps for PRL files from a REL file.
- WordTech's dBXL: Writing your own custom designed business program.
- Advanced CP/M: ZEX 5.0×The machine and the language.
- Programming for Performance: Assembly language techniques.
- Programming Input/Output With C: Keyboard and screen functions.
- The Z-System Corner: Remote access systems and BDS C.
- Real Computing: The NS320XX
- The Computer Corner.

Issue Number 41:

- Forth Column: ADTs, Object Oriented Concepts.
- Improving the Ampro LB: Overcoming the 88Mb hard drive limit.
- How to add Data Structures in Forth
- Advanced CP/M: CP/M is hacker's haven, and Z-System Command Scheduler.
- The Z-System Corner: Extended Multiple Command Line, and aliases.
- Programming disk and printer functions with C.
- LINKPRL: Making RSXes easy.
- SCOPY: Copying a series of unrelated files.
- The Computer Corner.

Issue Number 42:

- Dynamic Memory Allocation: Allocating memory at runtime with examples in Forth.
- Using BYE with NZCOM.
- C and the MS-DOS Screen Character Attributes.

· Forth Column: Lists and object oriented Forth.
· The Z-System Corner: Genie, BDS Z and Z-System Fundamentals.
· 68705 Embedded Controller Application: An example of a single-chip microcontroller application.
· Advanced CP/M: PluPerfect Writer and using BDS C with REL files.
· Real Computing: The NS 32000.
· The Computer Corner

The Computer Journal

Back Issues

Sales limited to supplies in stock.

Issue Number 43:

· Standardize Your Floppy Disk Drives.
· A New History Shell for ZSystem.
· Heath's HDOS, Then and Now.
· The ZSystem Corner: Software update service, and customizing NZCOM.
· Graphics Programming With C: Graphics routines for the IBM PC, and the Turbo C graphics library.
· Lazy Evaluation: End the evaluation as soon as the result is known.
· S-100: There's still life in the old bus.
· Advanced CP/M: Passing parameters, and complex error recovery.
· Real Computing: The NS32000.
· The Computer Corner.

Issue Number 44:

· Animation with Turbo C Part 1: The Basic Tools.
· Multitasking in Forth: New Micros F68FC11 and Max Forth.
· Mysteries of PC Floppy Disks Revealed: FM, MFM, and the twisted cable.
· DosDisk: MS-DOS disk format emulator for CP/M.
· Advanced CP/M: ZMATE and using lookup and dispatch for passing parameters.
· Real Computing: The NS32000.
· Forth Column: Handling Strings.
· Z-System Corner: MEX and telecommunications.
· The Computer Corner

Issue Number 45:

· Embedded Systems for the Tenderfoot: Getting started with the 8031.
· The Z-System Corner: Using scripts with MEX.
· The Z-System and Turbo Pascal: Patching TURBO.COM to access the Z-System.
· Embedded Applications: Designing a Z80 RS-232 communications gateway, part 1.
· Advanced CP/M: String searches and tuning Jetfind.
· Ahimation with Turbo C: Part 2, screen interactions.
· Real Computing: The NS32000.
· The Computer Corner.

Issue Number 46:

· Build a Long Distance Printer Driver.
Using the 8031's built-in UART for serial communications.
· Foundational Modules in Modula 2.
· The Z-System Corner: Patching The Word Plus spell checker, and the ZMATE macro text editor.
· Animation with Turbo C: Text in the graphics mode.
· Z80 Communications Gateway: Prototyping, Counter/Timers, and using the Z80 CTC.

Issue Number 47:

· Controlling Stepper Motors with the 68HC11F
· Z-System Corner: ZMATE Macro Language
· Using 8031 Interrupts
· T-1: What it is & Why You Need to Know
· ZCPR3 & Modula, Too
· Tips on Using LCDs: Interfacing to the 68HC705
· Real Computing: Debugging, NS32 Multitasking & Distributed Systems
· Long Distance Printer Driver: correction
· ROBO-SOG 90
· The Computer Corner

Issue Number 48:

· Fast Math Using Logarithms
· Forth and Forth Assembler
· Modula-2 and the TCAP
· Adding a Bernoulli Drive to a CP/M Computer (Building a SCSI Interface)
· Review of BDS "Z"
· PMATE/ZMATE Macros, Pt. 1
· Real Computing
· Z-System Corner: Patching MEX-Plus and TheWord, Using ZEX
· Z-Best Software
· The Computer Corner

Issue Number 49:

· Computer Network Power Protection
· Floppy Disk Alignment w/RTXEB, Pt. 1
· Motor Control with the F68HC11
· Controlling Home Heating & Lighting, Pt. 1
· Getting Started in Assembly Language
· LAN Basics
· PMATE/ZMATE Macros, Pt. 2
· Real Computing
· Z-System Corner
· Z-Best Software
· The Computer Corner

Issue Number 50:

· Offload a System CPU with the Z181
· Floppy Disk Alignment w/RTXEB, Pt. 2
· Motor Control with the F68HC11
· Modula-2 and the Command Line
· Controlling Home Heating & Lighting, Pt. 2
· Getting Started in Assembly Language Pt 2
· Local Area Networks
· Using the ZCPR3 IOP
· PMATE/ZMATE Macros, Pt. 3
· Z-System Corner, PCED
· Z-Best Software
· Real Computing, 32FX16, Caches
· The Computer Corner

Issue Number 51:

· Introducing the YASBEC
· Floppy Disk Alignment w/RTXEB, Pt 3
· High Speed Modems on Eight Bit Systems
· A Z8 Talker and Host
· Local Area Networks--Ethernet
· UNIX Connectivity on the Cheap
· PC Hard Disk Partition Table
· A Short Introduction to Forth
· Stepped Inference as a Technique for Intelligent Real-Time Embedded Control
· Real Computing, the 32CG160, Swordfish, DOS Command Processor
· PMATE/ZMATE Macros
· Z-System Corner, The Trenton Festival
· Z-Best Software, the Z3HELP System
· The Computer Corner

Issue Number 52:

· YASBEC, The Hardware
· An Arbitrary Waveform Generator, Pt. 1
· B.Y.O. Assembler...in Forth
· Getting Started in Assembly Language, Pt. 3
· The NZCOM IOP
· Servos and the F68HC11
· Z-System Corner, Programming for Compatibility
· Z-Best Software
· Real Computing, X10 Revisited
· PMATE/ZMATE Macros
· Controlling Home Heating & Lighting, Pt. 3
· The CPU280, A High Performance Single-Board Computer
· The Computer Corner

Issue Number 53:

· The CPU280
· Local Area Networks
· Am Arbitrary Waveform Generator
· Real Computing
· Zed Fest '91
· Z-System Corner
· Getting Started in Assembly Language
· The NZCOM IOP
· Z-BEST Software
· The Computer Corner

Issue Number 54:

· Z-System Corner
· B.Y.O. Assembler
· Local Area Networks
· Advanced CP/M
· ZCPR on a 16-Bit Intel Platform
· Real Computing
· Interrupts and the Z80
· 8 MHZ on a Ampro
· Hardware Heavenn
· What Zilog never told you about the Super8
· An Arbitary Waveform Generator
· The Development of TDOS
· The Computer Corner

Issue Number 55:

· Fuzzilogy 101
· The Cyclic Redundancy Check in Forth
· The Internetwork Protocol (IP)
· Z-System Corner
· Hardware Heaven
· Real Computing
· Remapping Disk Drives through the Virtual BIOS
· The Bumbling Mathmatician
· YASMEM
· Z-BEST Software
· The Computer Corner

Issue Number 56:

· TCJ - The Next Ten Years
· Input Expansion for 8031
· Connecting IDE Drives to 8-Bit Systems
· Real Computing
8 Queens in Forth
· Z-System Corner
· Kaypro-84 Direct File Transfers
· Analog Signal Generation
· The Computer Corner

Issue Number 57:

· Home Automation with X10
· File Transfer Protocols
· MDISK at 8 MHZ.
· Real Computing
· Shell Sort in Forth
· Z-System Corner
· Introduction to Forth
· DR. S-100
· Z AT Last!
· The Computer Corner

	U.S.	Foreign (Surface)	Foreign (Airmail)	Total
Subscriptions				
1year (6 issues)	$18.00	$24.00	$38.00	_____
2 years (12 issues)	$32.00	$44.00	$72.00	_____
Back Issues				
#20 thru #43	$3.50 ea.		$6.00 ea.	_____
#44 and up	$4.50 ea.		$7.00 ea.	_____
MicroC Kaypro Disks	$7.00 ea		$10.00 ea	_____

Items: _____

Subscription Total _____
Back Issues Total _____
MicroC Disks Total _____
California state Residents add 7.25% Sales TAX _____
Total Enclosed _____

Name: _____
Address: _____
_____
_____
My Interests: _____
_____

Payment is accepted by check or money order. Checks must be in US funds, drawn on a US bank. Personal checks within the US are welcome.

TCJ The Computer Journal

P.O. Box 535, Lincoln, CA 95648-0535
Phone (916) 645-1670

Regular Feature

Editorial Comment

UZI and Z180

Computer Corner

By Bill Kibler

Well another issues is done and ready for the press. It is late and I need to whip this out so I can finish the issue. Yet I actually do have items of importance (depending on your point of view) to comment on.

UZI on a Z180?

Rick Rodman's column this month got me thinking about the UZI program (Unix 7 for the Z80). The intro I included pointed out problems with such a limited memory space. That sort of ruled out doing much with it, until I considered the Z180. It would appear that a little re-working is in order, and major problems might be solved when it has a full Megabyte of memory for swapping instead of using the disk. Something we need to consider.

Speaking of the Z180, I have gotten some feed back on our proposed XT bus Z180 CPU card. So far the results is vary favorable. Everyone is in favor of it, except price seems to be a major concern. It would seem our competition is cheap clones that are falling into the $100 range and under. My position is still, that we are offering much more than you will get by buying that old clone.

What options, well 20MHZ CPU that would require at least a 40MHZ 386 to be as fast. Source code to operating systems and many utilities. Many times simpler design. Access to very cheap I/O cards. Ability to use LAN cards and other exotic I/O devices. Lots of boards that can be bought for $5 to play with and modify.

Of course you say, but why not the XT itself. My answer is go ahead and try. I spend many hours a week working on adapting PC/XT software and hardware for a living, and let me say without a doubt it is no fun! I hate anything that is written by Microsoft. The list of bugs in those products is endless. The problem is there is no way to work around most of those bugs. What is in store for the future, more complex problems and bugs.

So my position is the simpler the better. Z180 on a XT bus is simple and straight forward. I like that. Windows for Workgroups, however is another successful dog from Microsoft.

Windows for Workgroups

Well last week I set through four hours of seeing just how Windows for Workgroups works. It becomes very apparent that the main objective is killing any competition. The product contains something to compete with almost every vendor of LAN products. It contains mail programs, LAN interfaces, schedule programs, their new database products, and on into nausea.

The market is targeted at small businesses (2 to 5 users) with the idea of eliminating sneaker nets (that is when disk data is transferred in person from machine to machine). Yet the $25 network (1-800-628-7992 and tell them The Computer Journal sent you) can do the same thing with less overhead and no need for windows. I would say my productivity drop by 30% thanks to waiting for windows to do it's thing. Now do it all with a Z180 and a few serial ports and now we are talking simple and fast.

Z80 on a LAN?

One area I am looking into is tying CP/M systems together. Now most servers are MSDOS based (actually OS2 for LanMan or Netware for the other 70%.) And when it comes for raw storage the cost ratio for a newer PC based machine is definitely better than any other. So it only logically makes sense to use them for servers. Now to use a 386 with 8 Megabytes of memory for a word processor seems to be a little overkill to me.

A better alternative would be using a Z80 based machine, tuned for word processing and connected to the PC server. This seems a great ratio and balance of simple and complex ways to solve problems. The large and virtual nature of the server balances the speed and simplicity of the Z80 for handling characters on a screen.

I guess one area to consider is how about graphics on a Z80. When I look at how most of the more recent word processor are more like desktop publishers, I really don't see a problem. I even know of one magazine published solely on a PC based word processor. So in reality I find little that could not be done with a CP/M talking to a MSDOS server. So where is that old LAN hardware and software that started out on Z80s? There should still be some of the old code and products kicking around. Anybody got any? let me know will you.

Well I need to save room for the Kaypro disk catalog, so I hope you consider the Z180 on a PC bus and drop me or Herb a note about your ideas. Till later keep on computing.

K-1
MODEM PROGRAMS

1-DISK	DOC	3k	KMDM795	COM	14k	MODEMPAT	COM	12k
CRC	COM	3k	KMDM795	DOC	1k	SQ	COM	14k
CRC	DOC	1k	KMDM795	LIB	18k	SQ/USQ	DOC	2k
CRCKLIST	CRC	1k	MODEM7	COM	9k	TERM	DOC	2k
D	COM	3k	MODEM7	DOC	11k	TERM	MAC	4k
FLS	COM	9k	MODEM7+	COM	12k	USQ	COM	10k
KMDM795	AQM	50k	MODEMPAT	ASM	9k			

MODEMPAT sets up the SIO (serial port) for whatever baud rate, bits per character, stop bits, and parity you need.
MODEM7+ This is MODEM7 with the MODEMPAT already added so you can select the correct communications interface each time you enter MODEM7.
NOTE: See disk K28 for a later version of Modem 7. However, the disk K28 version limits you to 300 or 1200 baud. This Modem 7+ gives you the complete range from 300 to 9600 baud.
KMDM795 A superset of MODEM7, this program lets you set baud rate on the fly. However, it is not set up for changing the bits/char and such like MODEM7+.
TERM A disassembly of the TERM program that was distributed with the early KAYPRO IIs and 4s. With this commented source, you can configure it for your own needs.
SQUEEZE/UNSQUEEZE These programs squeeze and unsqueeze all kinds of files.

K-2
UTILITIES

2-DISK	DOC•	4k	DIRCHK	DOC	1k	MAST	CAT	1k
ALLOC	COM•	2k	DU-77	COM•	6k	QCAT	COM	1k
CAT	COM	1k	DU-77	DOC•	6k	SUPERSUB	COM	3k
CATALOG	DOC	8k	DUMPX	COM	3k	SUPERSUB	DOC	5k
COMPARE	COM•	2k	DUMPX	DOC	5k	UCAT	COM	2k
CRC	COM•	3k	FINDBD54	COM	2k	UNERA11	HLP	3k
CRC	DOC•	1k	FINDBD54	DOC	1k	UNERA19	COM	2k
CRCKLIST	CRC•	2k	FIX	COM	36k	UNERA19	DOC	3k
D	COM	3k	FIX	DOC	1k	UNLOAD	COM	1k
D	DOC	3k	FMAP	COM	4k	UNSPOOL	COM	2k
DASM	COM	10k	FORMFEED	COM	1k	UNSPOOL	DOC	14k
DASM	DOC	27k	FORMFEED	DOC	1k	WASH	COM•	3k
DIR-DUMP	COM•	1k	LISTT	COM	2k			
DIRCHK	COM	2k	LISTT	DOC	4k			

FORMFEED Run this program when you need a form feed on your printer.
DIR-DUMP Displays all the disk files and their locations, along with the user number.
D This super directory program produces an alphabetical list in vertical order.
WASH A file transfer/maintenance program. Forerunner (and subset) of SWEEP.
DUMPX Similar to the CP/M routine DUMP, DUMPX is incredibly more powerful and more useful (e.g. displays the ASCII equivalents).
SUPERSUB A definite improvement over CP/M's SUBMIT program.
ALLOC Produces a bit map of the disk and can write-protect files. Menu driven.
COMPARE Compares two files.
DU-77 Very powerful disk utility. The user interface isn't particularly easy to learn but you can do anything to a disk with this famous program.
LISTT Prints out CP/M files with headers, page numbers, and offsets from left margin.
DASM True Zilog format disassembler for 8080 and Z80 object (.COM) files.
UNSPOOL Lets you use your system and print at the same time (without one of those fancy print buffers).
FINDBD54 This super utility really checks a disk (without destroying the data on it) and reports bad sectors.
CATALOG A group of programs to create and maintain a directory of all the programs you have on all your disks.
UNERA Unerases files. Kinda handy.

MICRO CORNUCOPIA

CATALOG OF KAYPRO CP/M USER DISKS

Available in 5 1/4 190 K Disks, SSDD

- Dotted Programs will also run on Commodore 128s

TCJ The Computer Journal
P.O. Box 535, Lincoln, CA 95648-0535
Phone (916) 645-1670

K-3
GAMES

3-DISK	DOC•	3k	D	COM	3k	PRESSUP	COM•	8k
ADV-SAVE	DOC•	1k	GAMES	COM•	39k	WASH	COM•	3k
BIO	COM	14k	MAZE	COM	3k	WUMPUS	COM•	15k
CAVE0	•	1k	MM	C •	5k	WUMPUS	DOC•	4k
CAVE1	•	1k	MM	COM•	7k	WUMPUS	PAS•	12k
CAVE2	•	1k	OTHELLO	COM	22k	ZCHESS	COM•	8k
CAVE4	•	1k	OTHELLO	DOC	2k	ZCHESS	DOC•	3k
CAVE5	•	1k	PACMAN	COM•	18k			
CRC	COM•	3k	PACMAN	DOC•	2k			
CRCKLIST	CRC•	1k	PRESSUP	C •	8k			

PACMAN Almost the real thing! And on a system that has no graphics!
ZCHESS A real honest-to-gosh, competent (at least more competent than D) chess game with a 1 to 6 level look-ahead.
OTHELLO Othello is a game you learn in minutes, and master in years.
GAMES Seven games all put together as one program. You select from LIFE, HORSE RACES, MAZE, PATTERNS, BLACKJACK, ANIMAL, and GO-BANG.
WUMPUS Wumpus is a classic computer game. You have a choice of caves as you enter and you can map them as you go through.
PRESSUP Pressup is somewhat similar in play to Othello but is not as difficult to master.
MM This is the Master Mind game. You try to guess what characters the computer has chosen. The point of the game is to guess the characters in as few moves as possible.
BIO Generates biorhythm charts. This handy program does it all, complete with a graphic display of the results.
MAZE Generates random mazes (surprise!).

K-4
ADVENTURE

-KAYPRO	004•	0k	ADVI	DAT•	31k	ADVT	PTR•	15k
ADV	COM•	36k	ADVI	PTR•	4k			
ADVD	SAV•	0k	ADVT	DAT•	105k			

ADVENTURE Here it is! The latest, greatest, most cussed adventure ever devised by half-mortals. This cave is greatly expanded and the creatures are much smarter (smarter than I anyway). These files total 191K so there isn't room for the usual documentation. See ADV-SAVE.DOC on disk K-3 for documentation.

K-5
MX-80/GEMINI 10X GRAPHICS SOFTWARE

5-DISK	DOC	1k	GRAF	C	8k	HUGH	GRF	12k
BAR1	GRF	2k	GRAF	COM	46k	LINE	GRF	3k
BAR2	GRF	2k	GRAF	DOC	21k	MEDIUM	GRF	4k
CIRCLE	GRF	8k	GRAFCIRC	C	2k	SMALL	GRF	1k
CRC	COM	3k	GRAFFILE	C	2k	STAR	GRF	8k
CRCKLIST	CRC	1k	GRAFINVT	C	1k	TITLE	GRF	8k
D	COM	3k	GRAFLINE	C	2k			
DEMO	C	6k	GRAFPLOT	C	2k			
DEMO	COM	41k	GRAFUTIL	C	3k			

Don Brittain has created this graphics display package for use with MX-80, FX-80 and FX-100 with Graftrax, or Gemini 10X printers.
GRAF An interactive routine that allows you to create graphic images.
NOTE: You get the graphics on the printer only. Even the new Kaypros cannot display the kind of high-resolution graphics that this package can do. The .COM files will work with Epson MX-80 compatible printers only. If you are familiar with C and have Aztec C you can try modifying and recompiling the source to match a different printer.

K-6
TEXT UTILITIES

6-DISK	DOC•	4k	ERTRW4Y2	•	1k	ROFF2	C •	10k
CHOP	C •	1k	ENTAB	C •	2k	RTW	C •	1k
CHOP	COM•	4k	ENTAB	COM•	3k	RTW	COM•	4k
CP	C •	2k	EPSMODE4	ASM	7k	SIGNS	COM	10k
CP	COM•	4k	EPSMODE4	COM	2k	TRUNC	C •	1k
CRC	COM•	3k	ERRTW3T	•	1k	TRUNC	COM•	3k
CRCKLIST	CRC•	2k	FONT	DAT•	1k	WRAP	C •	7k
D	COM	3k	MS	C •	1k	WRAP	COM•	6k
DUMP	C •	4k	MS	COM•	4k	WRAP	DOC•.	3k
DUMP	COM•	5k	ROFF	COM•	18k	YUTTRYR	•	1k
EDIT	COM•	26k	ROFF	DOC•	8k			
EDIT	DOC•	10k	ROFF	H •	4k			
EDIT	HLP•	1k	ROFF1	C •	12k			

ROFF A UNIX-like text formatter.
SIGNS creates large block letters on your printer.
EPSMODE4 Set print modes on your Epson MX80.
EDIT Full-blown line editor similar UNIX's EX.
CHOP cuts off a file after n bytes. **CP** will copy from one file or device to another. **DUMP** outputs (to screen, disk, etc.) the HEX translation of a file. **MS** lets you specify how many linefeeds you want between lines (multiple space). **WRAP** is a simple formatter.
ENTAB replaces blanks with tabs wherever possible.
RTW removes trailing whitespace from a text file.
TRUNC truncates each line in a text file at a specified column.

K-7
SMALL C VERSION 2

7-DISK	DOC•	2k	ITOX	C •	1k	STD	H •	1k
ABS	C •	1k	LEFT	C •	1k	STDIO	H •	3k
CALL	ASM•	10k	LIB	H •	1k	STDIOA	H •	1k
CRC	COM•	3k	LIBASM	C •	1k	STRCMP	C •	1k
CRCKLIST	CRC•	2k	OUT	C •	1k	UTOI	C •	1k
DTOI	C •	1k	PRINTF	C •	2k	VM	H •	1k
HELLO	C •	1k	SAMPLE	DOC•	7k	VML	H •	1k
HELLO	HEX•	0k	SIGN	C •	1k	XTOI	C •	1k
IOLIB	ASM•	26k	SMALLC	DOC•	17k			
ITOD	C •	1k	SMALLC2	DOC•	6k			
ITOU	C •	1k	SMC	COM•	35k			

Here is an expanded version of Ron Cain's Small C. The additions to the language are substantial and include such important features as I/O redirection on STDIN and STDOUT, separate compilation of version 2 modules and many new library functions.

K-8
SOURCE OF SMALL C VERSION 2

ABS	C •	1k	CC33	C •	5k	LIBASM	C •	1k
CC	DEF•	4k	CC4	C •	1k	OUT	C •	1k
CC1	C •	6k	CC41	C •	7k	PRINTF	C •	2k
CC11	C •	11k	CC42	C •	9k	SIGN	C •	1k
CC12	C •	8k	CC80	DEF•	4k	SMCC	DEF•	4k
CC13	C •	8k	CRC	COM•	3k	STDIOA	H •	1k
CC2	C •	2k	CRCKLIST	CRC•	2k	STRCMP	C •	1k
CC21	C •	7k	DTOI	C •	1k	UTOI	C •	1k
CC22	C •	9k	ITOD	C •	1k	XTOI	C •	1k
CC3	C •	2k	ITOU	C •	1k			
CC31	C •	9k	ITOX	C •	1k			
CC32	C •	6k	LEFT	C •	1k			

K-8 contains the source of Small C version 2. You do not need this disk to use the compiler on K7.

K-9
GENERAL UTILITIES

9-DISK	DOC•	4k	EX14	COM	3k	INSTALL	SUB	1k
ALIENS	COM	14k	EX14	DOC	6k	PASSWORD	ASM•	8k
ALIENS	DOC	3k	EX14	SUB	5k	PASSWORD	COM•	1k
CRC	COM	3k	EX14	TST	1k	SNOOPY	TXT	3k
CRCKLIST	CRC	1k	FIND	ASM•	8k	SSED2	COM	11k
D	COM	3k	FIND	COM•	1k	TREK	COM•	27k
DIF/SSED	DOC	5k	FIND	DOC•	1k	ZCPR	DOC	6k
DIF2	COM	15k	FIX	COM	29k	ZCPR	HEX	6k
EX14	ASM	25k	FIX	DOC	1k			

EX14 Great replacement for both SUBMIT and XSUB.
FIX A fantastic disk utility that allows you do nearly anything on a disk.
FIND Searches the disk for a string of uppercase characters.
PASSWORD allows you to password-protect files.
ALIENS Space Invaders patched for the Kaypro.
TREK You finally get a chance to command the Starship Enterprise.
ZCPR This CCP replacement will look on drive A for a COM file when you're logged in on drive B, page during TYPE, and more.
DO NOT ATTEMPT TO INSTALL ZCPR ON THE KAYPRO 10!
DIF2/SSED With these utilities you can update someone else's copy of a program by simply creating a file of the updates (using DIF2) and then sending only the differences.

K-10
Z80 AND LINKING ASSEMBLERS

10-DISK	DOC	3k	CROWECPM	DOC	8k	LASM	DOC	5k
CRC	COM	3k	CROWECPM	PRN	0k	PRINTPRN	ASM	5k
CRCKLIST	CRC	1k	CROWECPM	Z80	143k	PRINTPRN	COM	1k
CROWECPM	COM	9k	LASM	COM	6k			

CROWECPM If you've been looking for a good, basic Z80 assembler for $8, this is it. We modified the CROWE assembler so that it would work on any CP/M system. Source code is included so you can extend it to your heart's content.
LASM A faster, linking, rewrite of the standard CP/M ASM assembler.
PRINTPRN This is a nice little program that makes it possible to list the .PRN files produced by the CROWECPM assembler on any printer.

K-11
CHECKBOOK PROGRAM - LIBRARY UTILITIES

11-DISK	DOC•	3k	EXAMPL	CHK•	1k	LU	COM•	18k
CHECKS	COM•	10k	EXAMPL	DOC•	6k	LU	DOC•	23k
CHECKS	DOC•	19k	EXAMPL	HAM•	2k	LUDEF1	DOC•	7k
CHECKS	Z80•	54k	LDIR	COM•	6k	PGLST	COM•	8k
CRC	COM•	3k	LDIR	DOC•	3k	PGLST	DOC•	1k
CRCKLIST	CRC•	1k	LDIR	LBR•	13k	VLIST	COM•	2k
DISPLAY	COM•	3k	LIBRARY	LBR•	1k	VLIST	DOC•	2k
DISPLAY	DOC•	3k	LBUN	COM•	2k			

DISPLAY is like the TYPE command except it allows you to page forward or backward.
VLIST uses the cursor keys to vary the scrolling speed of a file you are TYPEing.
GLST reformats a long text file of short words into columns.
CHECKS keeps track of which checks are tax deductable and which should be charged to various accounts. It also keeps a running balance.
LU (LIBRARY UTILITIES) This combination of utilities allows you to extract individual files or run COM files from the library without extracting them first. You can also build your own library files.

K-12
KAYPRO FORTH

12-DISK	DOC•	2k	FORTH	COM•	7k	KFORTH	COM•	16k
CRC	COM•	3k	FORTH	DOC•	12k			
CRCKLIST	CRC•	1k	FORTH	SCR•	150k			

There are two FORTHs on this disk. **FORTH** is true Fig FORTH. **KFORTH** has been extended and includes its own screen editor, decompiler, and 8080 assembler.

K-13
SOURCE FOR FIG-FORTH

13-DISK	DOC	1k	CRCKLIST	CRC	1k	FORTH	COM	7k
CRC	COM	3k	FORTH	ASM	46k			

This disk contains the source of the Fig FORTH on disk K12.

K-14
SMARTMODEM PROGRAMS

14-DISK	DOC	3k	MODEMPAT	ASM	9k	SMODEMK	MAC	83k
BYE	COM	2k	MODEMPAT	COM	12k	XMODEM	COM	4k
CRC	COM	3k	PHONE	001	1k	XMODEM	DOC	9k
CRCKLIST	CRC	1k	PHONE	BAK	1k			
D	COM	3k	SMODEM	DOC	24k			
KAYTERM	DOC	5k	SMODEMK	COM	10k			

SMODEMK Modem7 setup for the Kaypro and SmartModem.
XMODEM lets a remote user upload and download files.
BYE connects and disconnects your system from the phone line.

K-15
HARD DISK UTILITIES

15-DISK	DOC	6k	CRCKLIST	CRC	1k	MULTCOPY	COM	2k
B1	COM	1k	D	COM	3k	MULTCOPY	DOC	1k
BACKUP	ASM	31k	FLOPCOPY	ASM	17k	SWEEP	COM	28k
BACKUP	COM	4k	FLOPCOPY	COM	3k	SWEEP38	DOC	14k
BACKUP	DOC	7k	FLOPCOPY	DOC	2k	U	ASM	4k
BIGBURST	ASM	12k	MAKE	COM	3k	U	COM	1k
BIGBURST	COM	2k	MDIR	ASM	10k	UNSQ	COM	12k
BIGBURST	DOC	3k	MDIR	COM	1k			
CRC	COM	3k	MOVE	COM	1k			
CRC	DOC	1k	MULTCOPY	ASM	14k			

BACKUP is a lifesaver. It backs up your hard disk onto as many floppy disks as necessary.
FLOPCOPY makes it easy to back up floppy disks on a system with one floppy drive and a hard disk. It uses the hard disk as a buffer.
BIGBURST breaks apart any file that is too large to fit onto a floppy in one piece.
MULTCOPY operates like PIP, but will prompt you when the floppy is full so you can change floppies.
MDIR This directory program displays files under ALL user areas and alphabetically sorts the entries in each area.
MAKE/MOVE Two utilities for moving files between user areas.
U lets you change user area and drive with one command.
SWEEP lets you do nearly everything you normally do with PIP, only easier. It's like TYPE, ERA, DIR, and PIP in one program.

K-16
PASCAL COMPILER

16-DISK	DOC•	1k	HW5DATA	•	1k	PSTACK	DOC•	2k
COMPARE	COM•	2k	PASCAL	DOC•	4k	REGEN	DOC•	3k
CRC	COM•	3k	PASYNTAX	DOC•	6k	RTP	ASM•	11k
CRC	DOC•	1k	PC	SUB•	1k	RTP	COM•	2k
CRCKLIST	CRC•	2k	PFET	COM•	7k	STIRLING	COM•	2k
DISK	DOC•	3k	PFET	PAS•	11k	STIRLING	PAS•	1k
EQ	COM•	3k	PLAYDATA	•	1k	TESTER	PAS•	4k
EQ	PAS•	2k	PLAYKAL	PAS•	13k	VALIDATE	SUB•	1k
EX14	COM•	3k	POPS	DOC•	2k	XA	OCO•	1k
EX14	DOC•	6k	POWTWO	PAS•	1k	XA	PCO•	1k
FWD	PAS•	1k	PPC	COM•	16k			
HW5	COM•	6k	PPC	DOC•	12k			
HW5	PAS•	15k	PPC	PAS•	26k			

As it stands, this version of Pascal supports only a subset of the language, but since the disk also contains complete source for the compiler, you can extend it if you wish. The compiler is written in Pascal and it compiles itself.

K-17
Z80 TOOLS

17-DISK	DOC•	2k	DASM	DOC•	27k	XLATE2	DOC•	3k
CRC	COM•	3k	DASM	MAC•	74k	XLATE2	MAC•	37k
CRC	DOC•	1k	DASMZLG	MAC•	17k			
CRCKLIST	CRC•	1k	DISASMBL	DOC•	5k			
DASM	COM•	10k	XLATE2	COM•	5k			

XLATE2 translates 8080 assembly language programs into Z80 mnemonics.
DASM is a friendlier version of the disassembler on disk K2.

K-18
SYSTEM DIAGNOSTICS

10DSKTST	COM	3k	DISKALGN	MAC	9k	MDIAGTTY	COM	4k
18-DISK	DOC	2k	DISKTEST	DOC	11k	MDIAGTTY	MAC	36k
2DISKTST	COM	3k	DISKTEST	MAC	24k	MDIAGXXX	DOC	13k
4DISKTST	COM	3k	LOWLPT	COM	4k	MEMRS	ASM	21k
CRC	COM	3k	LOWTTY	COM	4k	MEMRS	COM	2k
CRCKLIST	CRC	1k	MDIAGLPT	COM	4k	MEMRS	DOC	7k
DISKALGN	COM	1k	MDIAGLPT	MAC	36k			

MEMRS performs a number of memory diagnostic tests from about 1000H to the bottom of BDOS (the TPA).
MDIAGTTY/MDIAGLPT Two programs to test from 1000H to the top of memory. The test procedures are different and a bit more exhaustive than those in MEMRS.
LOWTTY/LOWLPT Revisions of MDIAGTTY and MDIAGLPT that test the portion of memory that is not tested by the others (0000H to 1000H).
DISKTST Nice little disk diagnostic which doesn't require an alignment disk or scope.
DISKALGN simply positions your drive head to a specified track. This makes alignment with a standard analog alignment disk much easier (you still need an alignment disk, a scope and a manual).

K-19
PROWRITER GRAPHICS SOFTWARE

19-DISK	DOC	1k	GRAF	C	8k	HUGH	GRF	12k
BAR1	GRF	2k	GRAF	COM	41k	LINE	GRF	3k
BAR2	GRF	2k	GRAF	DOC	21k	MEDIUM	GRF	4k
CIRCLE	GRF	8k	GRAFCIRC	C	2k	SMALL	GRF	1k
CRC	COM	3k	GRAFFILE	C	2k	STAR	GRF	8k
CRCKLIST	CRC	1k	GRAFINVT	C	2k	TITLE	GRF	8k
D	COM	3k	GRAFLINE	C	2k			
DEMO	C	7k	GRAFPLOT	C	2k			
DEMO	COM	35k	GRAFUTIL	C	5k			

Don Brittain's Prowriter version of disk K5. See the documentation on that disk for more information.

K-20
PROGRAMS FOR MICROSPHERE'S COLOR GRAPHICS BOARD

20-DISK	DOC	2k	PACINIT	C	8k	PIE	DOC	2k
CRC	COM	3k	PACMAN	C	7k	PORTS	O	1k
CRC	DOC	1k	PACMAN	COM	24k	SKETCH	BAS	26k
CRCKLIST	CRC	1k	PACMAN	SCR	14k	SKETCH	COM	19k
HIGHSCR	PAC	1k	PACMONST	C	5k	SKETCH	DOC	6k
PAC		1k	PACUTIL	C	6k	SKETCH	SCR	14k
PAC	SUB	1k	PIE	BAS	9k	TEXT	OUT	10k
PACDEFS	H	3k	PIE	COM	26k			

SBASIC Pacman, sketching, and pie chart programs for the color graphics board.

K-21
SBASIC PROGRAMS AND SCREEN DUMP

ADDRESS	DAT	0k	DUMP84	MAC	8k	MSTRMIND	COM•	12k
CRC	COM•	3k	HANGMAN	COM	21k	SCREEN	BAS•	2k
CRCKLIST	CRC•	2k	HANGMAN	DOC	2k	SCREEN	COM•	8k
DIR+	COM•	4k	INSCREEN	BAS	6k	SCREEN	DOC•	1k
DIR+	DOC•	1k	INSCREEN	COM	11k	SCROLL	COM•	1k
DRIVER	BAS•	12k	INSCREEN	DOC	9k	SCROLL	DOC	2k
DRIVER	COM•	13k	MATH	BAS•	4k	WORD	COM	8k
DRIVER	DOC•	3k	MATH	COM•	9k	WORD	LIB	1k
DUMP24	COM	1k	MATH	DOC•	1k	XLATE	BAS•	2k
DUMP24	DOC	3k	MATH1	BAS•	3k	XLATE	COM•	5k
DUMP24	MAC	6k	MATH1	COM•	13k	XLATE	DOC•	3k
DUMP84	COM	1k	MATHMIND	DOC•	1k			

DIR+ This disk utility does everything PIP does and much more.
DRIVER creates a menu of .COM files on drives A and B and executes selected file.
DUMP24 Screen to printer dump for older (83) model Kaypros.
DUMP84 The same as the above program for the new generation Kaypros.
HANGMAN Traditional word guessing game in SBASIC.
INSCREEN SBASIC screen input program.
MATH Micro C's attempt at structured SBASIC coding using Math1 as the victim.
MATH1 The math game in its original unstructured form.
MSTRMIND Mastermind game written in SBASIC.
SCREEN displays all displayable screen characters.
SCROLL allows viewing of files by scrolling forward or backward.

	U.S.	Foreign (Surface)	Foreign (Airmail)	Total
Subscriptions				
1year (6 issues)	$18.00	$24.00	$38.00	_____
2 years (12 issues)	$32.00	$44.00	$72.00	_____
Back Issues				
#20 thru #43	$3.50 ea.		$6.00 ea.	_____
#44 and up	$4.50 ea.		$7.00 ea.	_____
MicroC Kaypro Disks	$7.00 ea		$10.00 ea	_____

Subscription Total _____
Back Issues Total _____
MicroC Disks Total _____
California state Residents add 7.25% Sales TAX _____
Total Enclosed _____

Name: _____
Address: _____
_____
_____
My Interests: _____
_____

Payment is accepted by check or money order. Checks must be in US funds, drawn on a US bank. Personal checks within the US are welcome.

TCJ The Computer Journal
P.O. Box 535, Lincoln, CA 95648-0535
Phone (916) 645-1670

Discover
The Z-Letter
The Z-letter is the only monthly publication for CP/M and Z-System. Eagle computers and Spellbinder support. Licensed CP/M distributor.

Subscriptions: $18 US, $22 Canada and Mexico, $36 Overseas. Write or call for free sample.

The Z-Letter
Lambda Software Publishing
720 South Second Street
San Jose CA 95112-5820
(408) 293-5176

Advent Kaypro Upgrades
TurboROM. Allows flexible configuration of your entire system, read/write additional formats and more, only $35.
Personality Decoder Boards
Run more than two drives when using TurboROM, $25.
Hard Drive Conversion Kits. Call or write for availability & pricing.

Call (916)483-0312
eves, weekends or write
Chuck Stafford
4000 Norris Ave.
Scaramento, CA 95824

TCJ MARKET PLACE
Advertising for small business
First Insertion: $50
Reinsertion: $35

Rates include typesetting.
Payment must accompany order.
VISA, MasterCard, Discover, Diner's Club, Carte Blanche, JCB, EuroCard accepted.
Checks, money orders must be US funds. Resetting of ad consitutes a new advertisement at first time insertion rates.
Mail ad or contact ,
The Computer Journal
P.O. Box 535
Lincoln, CA 95648-0535

CP/M SOFTWARE

100 page Public Domain Catalog, $8.50 plus $1.50 shipping and handling. New Digital Research CP/M 2.2 manual, $19.95 plus $3.00 shipping and handling. Also, MS/PC-DOS Software. Disk Copying, including AMSTRAD. Send self addressed, stamped envelope for free Flyer, Catalog $1.00

Elliam Associates
Box 2664
Atascadero, CA 93423
805-466-8440

8 BITS *and Change*

CLOSING OUT SALE!

All 12 Back Issues

for only $40

Send check to

Lee Bradley
24 East Cedar Street
Newington, CT 06111
(203) 666-3139 voice
(203) 665-1100 modem

S-100/IEEE-696

Compupro
Cromemco
IMSAI
and more!

Cards • Docs • Systems

Dr. S-100

Herb Johnson,
CN 5256 #105,
Princeton, NJ 08543
(609) 588-5316

SUPPORT OUR ADVERTISERS TELL THEM "I SAW IT IN TCJ"

Z80 STD USERS!

Cost Effective Upgrade
Clock Speeds to 10 MHz
1 Mbyte On-board Memory

Increase your system performance and reliability while reducing your costs by replacing three of the existing cards in your system with one **Superintegrated** Z80 Card from Zwick Systems.

A **Superintegrated** Card in your system protects your software investment, requiring only minor changes to your mature Z80 code. You can increase your processing performance by up to 300 percent in a matter of days!

Approximatly 35 percent of each **Superintegrated** Card has been reserved for custom I/O functions including A/D, D/A, Industrial I/O, Parallel Ports, Serial Ports, Fax and Data Modems or almost any other form of I/O that you are currently using.

Call or Fax today for complete information on this exciting new line of **Superintegrated** Cards and upgrade your system the easy way!

ZWICK SYSTEMS INC.
Tel (613) 726-1377, Fax (613) 726-1902